

“Advanced” Isn’t Always Better

*Assessing the Safety, Security, and Environmental
Impacts of Non-Light-Water Nuclear Reactors*

Edwin Lyman

March 2021

© 2021 Union of Concerned Scientists
All Rights Reserved

Edwin Lyman is the director of nuclear power safety in the UCS Climate and Energy Program.

The Union of Concerned Scientists puts rigorous, independent science to work to solve our planet's most pressing problems. Joining with people across the country, we combine technical analysis and effective advocacy to create innovative, practical solutions for a healthy, safe, and sustainable future.

This report is available online (in PDF format) at www.ucsusa.org/resources/advanced-isnt-always-better and <https://doi.org/10.47923/2021.14000>

Designed by:
David Gerratt, Acton, MA
www.NonprofitDesign.com

Cover photo: Argonne National Laboratory/Creative Commons (Flickr)

Printed on recycled paper.

[CONTENTS]

- vi Figures, Tables, and Boxes
- vii Acknowledgments

EXECUTIVE SUMMARY

- 2 Key Questions for Assessing NLWR Technologies
- 2 Non-Light Water Reactor Technologies
- 4 Evaluation Criteria
- 5 Assessments of NLWR Types
- 8 Safely Commercializing NLWRs: Timelines and Costs
- 9 The Future of the LWR
- 9 Conclusions of the Assessment
- 11 Recommendations
- 12 Endnotes

CHAPTER 1

- 13 **Nuclear Power: Present and Future**
- 13 Slower Growth, Cost and Safety Concerns
- 14 Can Non-Light-Water Reactors Revive Nuclear Power's Prospects?
- 15 A Note on Terminology
- 15 The Need to Fully Vet Claims About NLWRs
- 16 NLWRs: Past and Present
- 19 A Host of Challenges Even for More Mature NLWR Designs
- 22 Nuclear Power Growth and Climate Change Mitigation
- 22 Is Development of NLWRs Essential for Nuclear Power's Future?

CHAPTER 2

- 24 **Nuclear Power Basics**
- 24 Nuclear Chain Reactions in the Reactor Core
- 25 The Components of Nuclear Fuels
- 28 Thermal and Fast Reactors

- 29 Fuel Burnup and Refueling
- 30 Reactor Safety Considerations

CHAPTER 3

- 32 **Nuclear Power Sustainability**
- 32 Two Primary Goals for Increasing Sustainability
- 33 The Challenging and Conflicting Goals of Sustainability
- 34 High-Level Waste Reduction
- 39 Uranium Utilization Efficiency

CHAPTER 4

- 43 **Nuclear Proliferation and Terrorism Risks of Nuclear Power**
- 44 International Standards for Detecting Diversion of Nuclear Materials and Protecting Them from Theft
- 46 NLWRs: Enrichment Issues
- 47 Nuclear Terrorism and Proliferation Concerns of HALEU
- 49 Reprocessing and NLWRs

CHAPTER 5

- 55 **Liquid Sodium–Cooled Fast Reactors**
- 55 History and Current Status
- 57 Fast Reactors: Cost Considerations
- 59 Safety
- 66 Sustainability and Proliferation/Terrorism Risk
- 66 Sustainability of Once-Through Fast Reactors
- 67 Time Scale and Costs

CHAPTER 6

- 74 **High-Temperature Gas–Cooled Reactors**
- 74 The Technology
- 75 History and Current Status
- 77 Safety
- 81 Sustainability

- 83 Proliferation/Terrorism Risk
- 85 Readiness for Commercial Demonstration and Near-Term Deployment

CHAPTER 7

- 89 **Molten Salt Reactors**
- 90 History and Current Status
- 91 Safety
- 94 Sustainability
- 102 Proliferation/Terrorism Risk
- 104 Readiness for Commercial Demonstration and Deployment

CHAPTER 8

- 107 **Breed-and-Burn Reactors**
- 107 The Rationale for Breed-and-Burn Reactors
- 110 Breed-and-Burn Reactor Designs
- 112 Sustainability
- 113 Proliferation/Terrorism Risk

CHAPTER 9

- 115 **Conclusions and Recommendations**
- 115 Conclusions
- 116 Recommendations

APPENDIX

- 118 **Simple Models of Fast Burner/Breeder Cycles**
- 118 How Long Would It Take to Reduce Transuranics by a Factor of 10 with a Burner Reactor System?
- 120 How Long Would It Take for a Fast Breeder Reactor to Extract 100 Times As Much Energy from Uranium Ore as an LWR?
- 122 Uranium Utilization Efficiency of Burner Reactors vs. Breeder Reactors
- 122 Waste Reduction Factor of a Breeder Reactor

- 124 **Endnotes**

- 126 **References**

[FIGURES, TABLES, AND BOXES]

FIGURES

- 119 Figure A.1. Heavy Metal Mass Flow for Startup and Transition Core for a Fast Burner Reactor
- 119 Figure A-2. Annual Heavy Metal Mass Flow at Equilibrium for a Fast Burner Reactor
- 121 Figure A-3. Production of Initial Core and First Reload for a Fast Breeder Reactor
- 121 Figure A-4. Annual Heavy Metal Mass Flow at Equilibrium for a Fast Breeder Reactor

TABLES

- 4 Table ES-1. Current Status of US NLWR Projects
- 6 Table ES-2. How NLWRs Compare with LWRs on Safety, Sustainability, and Proliferation Risk
- 17 Table 1. Past and Present Demonstration Reactors Worldwide

BOXES

- 19 Box 1. Stages of Advanced Reactor Development
- 26 Box 2. Reactor Stability: Controlling the Chain Reaction
- 37 Box 3. Reprocessing and Recycling: Turning Spent Fuel into Fresh Fuel
- 38 Box 4. What Level of Transuranic Reduction in Radioactive Waste Would Make a Real Difference?
- 50 Box 5. The Ups and Downs of US Reprocessing Policy
- 61 Box 6. Reactivity Effects in Fast Reactors
- 71 Box 7. The Pyroprocessing Files
- 95 Box 8. Protactinium and the Thorium Fuel Cycle
- 99 Box 9. Transatomic Power: A Cautionary Tale

[ACKNOWLEDGMENTS]

This report was made possible with the support of the John D. and Catherine T. MacArthur Foundation and UCS members.

The author would like to thank the many UCS staff members, past and present, who provided invaluable input into this report: Angela Anderson, Rachel Cleetus, Rob Cowin, Peter Frumhoff, Lisbeth Gronlund, Kathleen Rest, and Steve Clemmer. He would like to thank Steve Fetter, Richard Garwin, Scott Kemp, and Frank von Hippel for reviewing the report and providing useful comments. He is very grateful to Cynthia DeRocco and Bryan Wadsworth for overseeing the report's production, and to Hannah Poor and Sital Sathia for their invaluable assistance. Finally, he would like to thank Karin Matchett, Marc S. Miller, and Heather Tuttle for their excellent editing and David Gerratt for his deft layout and design.

Organizational affiliations are listed for identification purposes only. The opinions expressed herein do not necessarily reflect those of the organizations that funded the work or the individuals who reviewed it. The Union of Concerned Scientists bears sole responsibility for the report's contents.

Executive Summary

The future of nuclear power is uncertain. Because nuclear power is a low-carbon way to generate electricity, there is considerable interest in expanding its role to help mitigate the threat of climate change. However, the technology has fundamental safety and security disadvantages compared with other low-carbon sources. Nuclear reactors and their associated facilities for fuel production and waste handling are vulnerable to catastrophic accidents and sabotage, and they can be misused to produce materials for nuclear weapons. The nuclear industry, policymakers, and regulators must address these shortcomings fully if the global use of nuclear power is to increase without posing unacceptable risks to public health, the environment, and international peace and security.

Despite renewed enthusiasm for nuclear power in many quarters, its recent growth has been far slower than anticipated 10 years ago. No doubt, the March 2011 Fukushima Daiichi accident in Japan, which resulted in three reactor meltdowns and widespread radiological contamination of the environment, has contributed to nuclear power's stagnation. Even more significant has been the high cost of building new reactors relative to other sources of electricity—primarily natural gas but also, increasingly, renewable energy sources such as wind and solar. The current rate of construction of new nuclear plants around the world barely outpaces the retirements of operating plants that reach the ends of their lifetimes or are no longer economic.

In the United States, new nuclear plants have proven prohibitively expensive and slow to build, discouraging private investment and contributing to public skepticism. In the 2000s, amid industry hopes of a nuclear renaissance,

the Nuclear Regulatory Commission (NRC) received applications to build more than two dozen new reactors. All were evolutionary versions of the light-water reactor (LWR), the type that comprises almost all operating reactors in the United States and most other countries with nuclear power. Companies such as Westinghouse, which developed the AP1000, promised these LWR variants could be built more quickly and cheaply while enhancing safety. But prospective purchasers cancelled nearly all of those proposals even before ground was broken, and the utilities that started building two AP1000 reactors at the V.C. Summer plant in South Carolina abandoned the project after it experienced significant cost overruns and delays. Only one project remains—two AP1000 units at the Alvin W. Vogtle plant in Georgia—but its cost has doubled, and construction is taking more than twice as long as originally estimated.

Almost all nuclear power reactors operating and under construction today are LWRs, so called because they use ordinary water (H₂O) to cool their hot, highly radioactive cores. Some observers believe that the LWR, the industry workhorse, has inherent flaws that are inhibiting nuclear power's growth. In addition to its high cost and long construction time, critics point to—among other things—the LWR's susceptibility to severe accidents (such as the meltdowns at Fukushima), their inefficient use of uranium, and the long-lived nuclear wastes they generate.

In response, the US Department of Energy's national laboratories, universities, and numerous private vendors—from large established companies to small startups—are pursuing the development of reactors that differ fundamentally from LWRs. These non-light-water reactors

(NLWRs) are cooled not by water but by other substances, such as liquid sodium, helium gas, or even molten salts.¹

NLWRs are sometimes referred to as “advanced reactors.” However, that is a misnomer for most designs being pursued today, which largely descend from those proposed many decades ago. At least one NLWR concept, the liquid metal-cooled fast reactor, even predates the LWR. Nevertheless, NLWR designers claim such reactors have innovative features that could disrupt the nuclear power industry and solve its problems. They state variously that their designs could lower costs, be built quickly, reduce the accumulation of nuclear waste, use uranium more efficiently, improve safety, and reduce the risk of nuclear proliferation. More specifically, they cite the advantages of features such as passive shutdown and cooling, the ability to consume or recycle nuclear waste, and the provision of high-temperature process heat for industrial applications such as hydrogen production. And some NLWR vendors claim that they can demonstrate, license, and deploy their designs within a decade or two.

Are these claims justified? How can we identify genuine innovations and recognize those that are likely unattainable? As with any technology, an independent reality check is needed. From self-driving cars to cheap flights to Mars, the Silicon Valley-style disruptive technology model of rapid, revolutionary progress is not always readily adaptable to other engineering disciplines. And nuclear energy, which requires painstaking, time-consuming, and resource-intensive research and development (R&D), is proving to be one of the harder technologies to disrupt.

In part, the nuclear industry’s push to commercialize NLWRs is driven by its desire to show the public and policymakers that there is a high-tech alternative to the static, LWR-dominated status quo: a new generation of “advanced” reactors. But a fundamental question remains: *Is different actually better?* The short answer is no. Nearly all of the NLWRs currently on the drawing board fail to provide significant enough improvements over LWRs to justify their considerable risks.

Key Questions for Assessing NLWR Technologies

It is critical that policymakers, regulators, and private investors fully vet the claims that the developers of NLWRs are making and accurately assess the prospects for both successful development *and* safe, secure, and cost-effective deployment. Given the urgency of the climate crisis, rigorous evaluation of these technologies will help our nation and others avoid wasting time or resources in the pursuit of

Given the urgency of the climate crisis, rigorous evaluation is needed to avoid wasting time or resources in the pursuit of high-risk energy concepts.

high-risk concepts that would be only slightly better—or perhaps worse—than LWRs.

Key questions to consider are the following:

- What are the benefits and risks of NLWRs and their fuel cycles compared with those of LWRs?
- Do the likely overall benefits of NLWRs outweigh the risks and justify the substantial public and private investments needed to commercialize them?
- Can NLWRs be safely and securely commercialized in time to contribute significantly to averting the climate crisis?

To help inform policy decisions on these questions, the Union of Concerned Scientists (UCS) has evaluated certain claims about the principal types of NLWRs. In particular, this report compares several classes of NLWRs to LWRs with regard to safety and security, the risks of nuclear proliferation and nuclear terrorism, and “sustainability”—a term that in this context includes the often-claimed ability of some NLWRs to “recycle” nuclear waste and use mined uranium more efficiently. The report also considers the potential for certain NLWRs to operate in a once-through, “breed-and-burn” mode that would, in theory, make them more uranium-efficient without the need to recycle nuclear waste—a dangerous process that has significant nuclear proliferation and terrorism risks.

Non-Light-Water Reactor Technologies

UCS considered these principal classes of NLWRs:

Sodium-cooled fast reactors (SFRs): These reactors are known as “fast reactors” because, unlike LWRs or other reactors that use lower-energy (or “thermal”) neutrons, the liquid sodium coolant does not moderate (slow down) the high-energy (or “fast”) neutrons produced when nuclear fuel undergoes fission. The characteristics and design features of these reactors differ significantly from those of LWRs, stemming from the properties of fast neutrons and the chemical nature of liquid sodium.

High-temperature gas-cooled reactors (HTGRs):

These reactors are cooled by a pressurized gas such as helium and operate at temperatures up to 800°C, compared with around 300°C for LWRs. HTGR designers developed a special fuel called TRISO (tristructural isotropic) to withstand this high operating temperature. HTGRs typically contain graphite as a moderator to slow down neutrons. There are two main variants of HTGR. A prismatic-block HTGR uses conventional nuclear fuel elements that are stationary; in a pebble-bed HTGR, moving fuel elements circulate continuously through the reactor core.

Molten salt-fueled reactors (MSRs): In contrast to conventional reactors that use fuel in a solid form, these use liquid fuel dissolved in a molten salt at a temperature of at least 650°C. The fuel, which is pumped through the reactor, also serves as the coolant. MSRs can be either thermal reactors that use a moderator such as graphite or fast reactors without a moderator. All MSRs chemically treat the fuel to varying extents while the reactor operates to remove radioactive isotopes that affect reactor performance. Therefore, unlike other reactors, MSRs generally require on-site chemical plants to process their fuel. MSRs also need elaborate systems to capture and treat large volumes of highly radioactive gaseous byproducts.

THE FUELS FOR NON-LIGHT-WATER REACTORS

Today's LWRs use uranium-based nuclear fuel containing less than 5 percent of the isotope uranium-235. This fuel is produced from natural (mined) uranium, which has a uranium-235 content of less than 1 percent, in a complex industrial process called uranium enrichment. Fuel enriched to less than 20 percent U-235 is called "low-enriched uranium" (LEU). Experts consider it a far less attractive material for nuclear weapons development than "highly enriched uranium" (HEU), with a U-235 content of at least 20 percent.

The fuel for most NLWRs differs from that of LWRs. Some proposed NLWRs would use LEU enriched to between 10 and 20 percent uranium-235; this is known as "high-assay low enriched uranium" (HALEU).² While HALEU is considered impractical for direct use in a nuclear weapon, it is more attractive for nuclear weapons development than the LEU used in LWRs. Other types of NLWRs would use plutonium separated from spent nuclear fuel through a chemical process called reprocessing. Still others would utilize the isotope uranium-233 obtained by irradiating the element thorium. Both plutonium and uranium-233 are highly attractive for use in nuclear weapons.

Typically, the chemical forms of NLWR fuels also differ from those of conventional LWR fuel, which is a ceramic

material composed of uranium oxide. Fast reactors can use oxides, but they can also use fuels made of metal alloys or chemical compounds such as nitrides. The TRISO fuel in HTGRs consists of tiny kernels of uranium oxide (or other uranium compounds) surrounded by several layers of carbon-based materials. MSR fuels are complex mixtures of fluoride or chloride salt compounds.

The deployment of NLWRs also would require new industrial facilities and other infrastructure to produce and transport their different types of fuel, as well as to manage spent fuel and other nuclear wastes. These facilities may use new technologies that themselves would require significant R&D. They also may present different risks related to safety, security, and nuclear proliferation than do LWR fuel cycle facilities—important considerations for evaluating the whole system.

NON-LIGHT-WATER REACTORS: PAST AND PRESENT

In the mid-20th century, the Atomic Energy Commission (AEC)—the predecessor of today's Department of Energy (DOE) and the NRC—devoted considerable time and resources to developing a variety of NLWR technologies, supporting demonstration plants at various scales at sites around the United States. Owners of several of these reactors abandoned them after the reactors experienced operational problems (for example, the Fort St. Vrain HTGR in Colorado) or even serious accidents (the Fermi-1 SFR in Michigan).

Despite these negative experiences, the DOE continued R&D on various types of NLWR and their fuel cycles. In the 1990s, the DOE initiated the Generation IV program, with the goal of "developing and demonstrating advanced nuclear energy systems that meet future needs for safe, sustainable, environmentally responsible, economical, proliferation-resistant, and physically secure energy." Although Generation IV identified six families of advanced reactor technology, the DOE has given most of its subsequent support to SFRs and HTGRs.

Today, a number of NLWR projects at various stages of development are under way, funded by both public and private sources (Table ES-1, p. 4). With support from Congress, the DOE is pursuing several new NLWR test and demonstration reactors. It is proceeding with the design and construction of the Versatile Test Reactor (VTR), an SFR that it hopes to begin operating in the 2026–2031 timeframe. The VTR would not generate electricity but would be used to test fuels and materials for developing other reactors. In October 2020, the DOE selected two NLWR designs for demonstrating commercial power generation by 2027: the Xe-100, a small pebble-bed HTGR that would generate about 76 megawatts

TABLE ES-1. Current Status of US NLWR Projects

Reactor Type	Power Level	Developer	Funding	NRC Licensing Status	Planned Startup Date
Sodium-Cooled Fast Reactors					
Versatile Test Reactor	300 MWth ^a	DOE	DOE	Not NRC licensed	2026–2031
Natrium	840 MWth/345 MWe	TerraPower-GE Hitachi	50–50 cost share with /ARDP ^b	Preapplication	2025–2027
Aurora Powerhouse ^c	4 MWth/1.5 MWe	Oklo, Inc.	Mostly private; some DOE subsidy	Combined operating license accepted for technical review June 2020	Early 2020s
High-Temperature Gas-Cooled Reactors					
Xe-100	4 x 200 MWth (76–80 MWe)	X-Energy	50–50 cost share with ARDP	Preapplication	2025–2027
Molten Salt Reactors					
IMSR	440 MWth/ up to 195 MWe	Terrestrial Energy	Private	Preapplication	—
Hermes reduced-scale test reactor ^d	Full-scale Kairos reactor 320 MWth/ 140 MWe; reduced scale > 50 MWth	Kairos Power	80 percent ARDP; 20 percent private	Preapplication	2027

a MWth: megawatts of thermal energy. MWe: megawatts of electricity.
 b ARDP: DOE Advanced Reactor Demonstration Program
 c The Aurora is potassium-cooled, with liquid sodium bonding contained in the fuel rods.
 d The Hermes is not molten salt-fueled but uses TRISO fuel and a molten-salt coolant.

of electricity (MWe), and the 345 MWe Natrium, an SFR that is essentially a larger version of the VTR with a power production unit. The DOE is also providing funding for two smaller-scale projects to demonstrate molten salt technologies. In addition, the DOE, the Department of Defense (DOD), and a private company, Oklo, Inc., are pursuing demonstrations of so-called micro-reactors—very small NLWRs with capacities from 1 MWe to 20 MWe—and project that these will begin operating in the next few years. A number of universities also have expressed interest in building small NLWRs for research.

Congress would need to provide sufficient and sustained funding for any of these projects to come to fruition. This is far from assured—for example, funding for the VTR to date has fallen far short of what the DOE has requested, all but guaranteeing the project will be delayed.

THE GOALS OF NEW NUCLEAR REACTOR DEVELOPMENT

If nuclear power is to play an expanded global role to help mitigate climate change, new reactor designs should be demonstrably safer and more secure—and more economical—than the existing reactor fleet. Today’s LWRs remain far too vulnerable to Fukushima-like accidents, and the uranium

enrichment plants that provide their LEU fuel can be misused to produce HEU for nuclear weapons. However, developing new designs that are clearly superior to LWRs overall is a formidable challenge, as improvements in one respect can create or exacerbate problems in others. For example, increasing the physical size of a reactor core while keeping its power generation rate constant could make the reactor easier to cool in an accident, but it could also increase cost.

Moreover, the problems of nuclear power cannot be fixed through better reactor design alone. Also critical is the regulatory framework governing the licensing, construction, and operation of nuclear plants and their associated fuel cycle infrastructure. Inadequate licensing standards and oversight activities can compromise the safety of improved designs. A key consideration is the extent to which regulators require extra levels of safety—known as “defense-in-depth”—to compensate for uncertainties in new reactor designs for which there is little or no operating experience.

Evaluation Criteria

UCS has considered three broad criteria for assessing the relative merits of NLWRs and LWRs: safety and security,

sustainability, and risks associated with nuclear proliferation and nuclear terrorism.

One characteristic that UCS did not consider here is the ability of reactors to provide high-temperature process heat for industrial applications—sometimes cited as a major advantage of NLWRs. However, potential industrial users have demonstrated little interest in these applications to date, and will likely continue to be wary of co-locating nuclear power plants at their facilities until outstanding safety, security, and reliability issues are fully addressed. It is also doubtful that industrial users would want to assume the cost and responsibility of managing the reactors' nuclear wastes. Consequently, UCS regards the generation of high-temperature process heat as a secondary objective that would first require significant improvements in nuclear safety and security.³

Safety and security risk is the vulnerability of reactors and fuel cycle facilities to severe accidents or terrorist attacks that result in significant releases of radioactivity to the environment. Routine radioactive emissions are also a consideration for some designs. The UCS assessment primarily used qualitative judgments to compare the safety of reactor types, because quantitative safety studies for NLWRs with the same degree of accuracy and rigor as for LWRs are not yet available. Far fewer data are available to validate safety studies of NLWRs than of LWRs, which have accumulated a vast amount of operating experience.

Sustainability, in this context, refers to the amount of nuclear waste generated by reactors and fuel facilities that requires secure, long-term disposal, as well as to the efficiency of using natural (mined) uranium and thorium. Sustainability criteria can be quantified but typically have large uncertainties. To account for those uncertainties, this report considers that sustainability parameters, such as the amount of heat-bearing transuranic (TRU) elements requiring long-term geologic disposal, would have to improve by a factor of 10 or more to be significant.

Nuclear proliferation and nuclear terrorism risk is the danger that nations or terrorist groups could illicitly obtain nuclear-weapon-usable materials from reactors or fuel cycle facilities. LWRs operating on a once-through fuel cycle present relatively low proliferation and terrorism risks. However, any nuclear fuel cycle that utilizes reprocessing and recycling of spent fuel poses significantly greater nuclear proliferation and terrorism risks than do LWRs without reprocessing, because it provides far greater opportunities for diversion or theft of plutonium and other nuclear-weapon-usable materials. International safeguards and security measures for reactors and fuel cycles with reprocessing are costly and cumbersome, and they cannot fully compensate for the increased vulnerability resulting from separating

weapon-usable materials. Also using HALEU instead of less-enriched forms of LEU would increase proliferation and terrorism risks, although to a far lesser extent than using plutonium or uranium-233.

Nuclear proliferation is not a risk in the United States simply because it already possesses nuclear weapons and is designated as a nuclear-weapon state under the Nuclear Non-Proliferation Treaty. As such, it is not obligated to submit its nuclear facilities and materials for verification by the International Atomic Energy Agency (IAEA), although it can do so voluntarily. However, US reactor development does have implications for proliferation, both because US vendors seek to export new reactors to other countries and because other countries are likely to emulate the US program. The United States has the responsibility to set a good international example by ensuring its own nuclear enterprise meets the highest nonproliferation standards.⁴

Not all these criteria are of equal weight. UCS maintains that increasing safety and reducing the risk of proliferation and terrorism should take priority over increasing sustainability for new reactor development at the present time. Given that uranium is now cheap and abundant, there is no urgent need to develop reactors that use less. Even so, there would be benefits from reducing the need for uranium mining, which is hazardous to workers and the environment and historically has had a severe impact on disadvantaged communities. Developing more efficient reactors may become more useful if the cost of mined uranium increases significantly, whether due to resource depletion or strengthened protections for occupational health and the environment.

UCS also did not consider the potential for NLWRs to be more economical than LWRs. Although economics is a critical consideration and is interrelated with the criteria listed above, such an evaluation would depend on many open and highly uncertain issues, such as final design details, future regulatory requirements, and supply chain availability.

Assessments of NLWR Types

UCS has reviewed hundreds of documents in the available literature to assess the comparative risks and benefits of the three major categories of NLWR with respect to the three evaluation criteria (Table ES-2, p. 6).

SODIUM-COOLED FAST REACTORS

Safety and Security Risk: SFRs have numerous safety problems that are not issues for LWRs. Sodium coolant can burn if exposed to air or water, and an SFR can experience rapid power increases that may be hard to control. It is even possible that an SFR core could explode like a small nuclear bomb under

severe accident conditions. Of particular concern is the potential for a runaway power excursion: if the fuel overheats and the sodium coolant boils, an SFR's power will typically increase rapidly rather than decrease, resulting in a positive feedback loop that could cause core damage if not quickly controlled.

Chernobyl Unit 4 in the former Soviet Union, although not a fast reactor, had a similar design flaw—known as a “positive void coefficient.” It was a major reason for the reactor’s catastrophic explosion in 1986. A positive void coefficient is decidedly not a passive safety feature—and it cannot be fully eliminated by design in commercial-scale SFRs. To mitigate these and other risks, fast reactors should have additional engineered safety systems that LWRs do not need, which increases capital cost.

Sustainability: Because of the properties of fast neutrons, fast reactors do offer, in theory, the potential to be more sustainable than LWRs by either using uranium more efficiently or reducing the quantity of TRU elements present in the reactor and its fuel cycle. This is the only clear advantage of fast reactors compared with LWRs. However, once-through fast reactors such as the Natrium being developed by TerraPower, a company founded and supported by Bill Gates, would be less uranium-efficient than LWRs. To significantly increase sustainability, most fast reactors would require spent fuel reprocessing and recycling, and the reactors and associated fuel cycle facilities would need to operate continuously at extremely high levels of performance for many hundreds or even thousands of years. Neither government nor industry can guarantee that future generations will continue to operate

TABLE ES-2. How NLWRs Compare with LWRs on Safety, Sustainability, and Proliferation Risk

NLWR Types	Safety	Sustainability		Nuclear Proliferation/ Terrorism
		Long-Lived Waste Generation	Resource Efficiency	
Sodium-Cooled Fast Reactors				
Conventional burner or breeder (Plutonium/TRU, with reprocessing)	---	++	+	---
Conventional: Natrium (HALEU, once-through)	---	--	--	--
Breed-and-burn mode (HALEU, once-through)	---	--	++	+
High-Temperature Gas-Cooled Reactors				
Prismatic-block (HALEU, once-through)	N	-	-	-
Pebble-bed: Xe-100 (HALEU, once-through)	N	-	-	--
Molten Salt-Fueled Reactors				
Thermal: IMSR/TAP (LEU <5% U-235)	---	+	-	-
Thermal: Thorcon (HALEU/Thorium/U-233)	---	-	+	--
Thermal: Molten Salt Breeder (HALEU/Thorium/U-233)	---	++	++	---
Molten Salt Fast Reactor (TRU/Thorium/U-233)	---	+++	++	---

■ Significantly Worse
 ■ Moderately Worse
 ■ Slightly Worse
 ■ Not Enough Information
■ Slightly Better
■ Moderately Better
■ Significantly Better

and replace these facilities indefinitely. The enormous capital investment needed today to build such a system would only result in minor sustainability benefits over a reasonable timeframe.

Nuclear Proliferation/Terrorism: Historically, fast reactors have required plutonium or HEU-based fuels, both of which could be readily used in nuclear weapons and therefore entail unacceptable risks of nuclear proliferation and nuclear terrorism. Some SFR concepts being developed today utilize HALEU instead of plutonium and could operate on a once-through cycle. These reactors would pose lower proliferation and security risks than would plutonium-fueled fast reactors with reprocessing, but they would have many of the same safety risks as other SFRs. And, as pointed out, most once-through SFRs would actually be less sustainable than LWRs and thus unable to realize the SFR's main benefit. For this reason, these once-through SFRs are likely to be “gateway” reactors that would eventually transition to SFRs with reprocessing and recycling. The only exceptions—if technically feasible—are once-through fast reactors operating in breed-and-burn mode. However, the only breed-and-burn reactor that has undergone significant R&D, TerraPower’s “traveling-wave reactor,” was recently suspended after more than a decade of work, suggesting that its technical challenges proved too great.

HIGH-TEMPERATURE GAS-COOLED REACTORS

Safety and Security Risk: HTGRs have some attractive safety features but also a number of drawbacks. Their safety is rooted in the integrity of TRISO fuel, which has been designed to function at the high normal operating temperature of an HTGR (up to 800°C) and can retain radioactive fission products up to about 1,600°C if a loss-of-coolant accident occurs. However, if the fuel heats up above that temperature—as it could in the Xe-100—its release of fission products speeds up significantly. So, while TRISO has some safety benefits, the fuel is far from meltdown-proof, as some claim. Indeed, a recent TRISO fuel irradiation test in the Advanced Test Reactor in Idaho had to be terminated prematurely when the fuel began to release fission products at a rate high enough to challenge off-site radiation dose limits.

The performance of TRISO fuel also depends critically on the ability to consistently manufacture fuel to exacting specifications, which has not been demonstrated. HTGRs are also vulnerable to accidents in which air or water leaks into the reactor; this is much less of a concern for LWRs. And the moving fuel in pebble-bed HTGRs introduces novel safety issues.

Despite these unknowns, HTGRs are being designed without the conventional leak-tight containments that LWRs have—potentially cancelling out any inherent safety benefits

provided by the design and fuel. Given the uncertainties, much more testing and analysis are necessary to determine conclusively if HTGRs would be significantly safer than LWRs.

Sustainability: HTGRs are less sustainable than LWRs overall. They use uranium no more efficiently due to their use of HALEU, and they generate a much larger volume of highly radioactive waste. Although pebble-bed HTGRs are somewhat more flexible and uranium-efficient than prismatic-block HTGRs, the difference is not enough to overcome the penalty from using HALEU fuel.

Nuclear Proliferation/Terrorism: HTGRs raise additional proliferation issues compared with LWRs. Current HTGR designs use HALEU, which poses a greater security risk than the LEU grade used by LWRs, and TRISO fuel fabrication is more challenging to monitor than LWR fuel fabrication. Also, it is difficult to accurately account for nuclear material at pebble-bed HTGRs because fuel is continually fed into and removed from the reactor as it operates. On the other hand, it may be more difficult for a proliferator to reprocess TRISO spent fuel than LWR spent fuel to extract fissile material because the required chemical processes are less mature.

MOLTEN SALT-FUELED REACTORS

Safety and Security Risk: MSR advocates point to the fact that this type of reactor cannot melt down—the fuel is already molten. However, this simplistic argument belies the fact that MSR fuels pose unique safety issues. Not only is the hot liquid fuel highly corrosive, but it is also difficult to model its complex behavior as it flows through a reactor system. If cooling is interrupted, the fuel can heat up and destroy an MSR in a matter of minutes. Perhaps the most serious safety flaw is that, in contrast to solid-fueled reactors, MSRs routinely release large quantities of gaseous fission products, which must be trapped and stored. Some released gases quickly decay into troublesome radionuclides such as cesium-137—the highly radioactive isotope that caused persistent and extensive environmental contamination following the Chernobyl and Fukushima nuclear accidents.

Sustainability: A main argument for MSRs is that they are more flexible and can operate more sustainably than reactors using solid fuel. In theory, some MSRs would be able to use natural resources more efficiently than LWRs and generate lower amounts of long-lived nuclear waste. However, the actual sustainability improvements for a range of thermal and fast MSR designs are too small, even with optimistic performance assumptions, to justify their high safety and security risks.

Nuclear Proliferation/Terrorism: MSRs present unique challenges for nuclear security because it would be

very difficult to account for nuclear material accurately as the liquid fuel flows through the reactor. In addition, some designs require on-site, continuously operating fuel reprocessing plants that could provide additional pathways for diverting or stealing nuclear-weapon-usable materials.

MSRs could also endanger global nuclear security by interfering with the worldwide network of radionuclide monitors put into place to verify compliance with the Comprehensive Nuclear Test Ban Treaty after it enters into force.⁵ MSRs release vast quantities of the same radioactive xenon isotopes that are signatures of clandestine nuclear explosions—an issue that MSR developers do not appear to have addressed. It is unclear whether it would be feasible or affordable to trap and store these isotopes at MSRs to the degree necessary to avoid degrading the effectiveness of the monitoring system to detect treaty violations.

Safely Commercializing NLWRs: Timelines and Costs

Can NLWRs be deployed quickly enough to play a significant role in reducing carbon emissions and avoiding the worst effects of climate change? The 2018 special report of the UN's Intergovernmental Panel on Climate Change identified 85 energy supply pathways to 2050 capable of achieving the Paris Agreement target of limiting global mean temperature rise to 1.5°C. The median capacity of nuclear power in 2050 across those pathways is about 150 percent over the 2020 level. Taking into account planned retirements, this corresponds to the equivalent of at least two dozen 1,000 MWe reactors coming online globally each year between now and 2050—five times the recent global rate of new LWR construction. If the world must wait decades for NLWRs to be commercially available, they would have to be built even faster to fill the gap by 2050.

Some developers of NLWRs say that they will be able to meet this challenge by deploying their reactors commercially as soon as the late 2020s. However, such aggressive timelines are inconsistent with the recent experience of new reactors such as the Westinghouse AP1000, an evolutionary LWR. Although the AP1000 has some novel features, its designers leveraged many decades of LWR operating data. Even so, it took more than 30 years of research, development, and construction before the first AP1000—the Sanmen Unit 1 reactor in China—began to produce power in 2018.

How, then, could less-mature NLWR reactors be commercialized so much faster than the AP1000? At a minimum, commercial deployment in the 2020s would require bypassing two developmental stages that are critical for assuring safety and reliability: the demonstration of prototype reactors

Commercial deployment in the 2020s would require bypassing prototype stages that are critical for assuring safety and reliability.

at reduced scale and at full scale. Prototype reactors are typically needed for demonstrating performance and conducting safety and fuel testing to address knowledge gaps in new reactor designs. Prototypes also may have additional safety features and instrumentation not included in the basic design, as well as limits on operation that would not apply to commercial units.

In a 2017 report, the DOE asserted that SFRs and HTGRs were mature enough for commercial demonstrations without the need for additional prototype testing. For either of these types, the DOE estimated it would cost approximately \$4 billion and take 13 to 15 years to complete a first commercial demonstration unit, assuming that reactor construction and startup testing take seven years. After five years of operating the demonstration unit, additional commercial units could follow in the mid-2030s.

In contrast, for MSRs and other lower-maturity designs, the DOE report judged that both reduced-scale and full-scale prototypes (which the report referred to as “engineering” and “performance” demonstrations, respectively) would be needed before a commercial demonstration reactor could be built. These additional stages could add \$2 billion to \$4 billion to the cost and 20 years to the development timeline. The subsequent commercial demonstration would not begin until 2040; reactors would not be available for sale until the mid-2040s or even the 2050s.

In May 2020, after receiving \$160 million in initial congressional funding for the new Advanced Reactor Demonstration Program (ARDP), the DOE issued a solicitation for two “advanced” commercial demonstration reactors. In October 2020, the DOE chose SFR and HTGR designs—as one might expect given its 2017 technology assessment. The DOE estimates that these projects will cost up to \$3.2 billion each (with the vendors contributing 50 percent) for the reactors and their supporting fuel facilities. The department is requiring that the reactors be operational within seven years, a timeline—including NRC licensing, construction, fuel production, and startup testing—that it acknowledges is very aggressive.

However, even if this deadline can be met and the reactors work reliably, subsequent commercial units likely would not be ordered before the early 2030s. Moreover, it is far from certain that the two designs the DOE selected for the ARDP are mature enough for commercial demonstration. Past demonstrations of both SFRs and HTGRs have encountered safety and reliability problems. Additionally, for both reactor types, the DOE has chosen designs that differ significantly from past demonstration reactors.

In the 1990s, the NRC concluded that it would require information from representative prototype testing prior to licensing either of these reactor types—but no prototypes were ever built. More recently, in a letter to the NRC, the agency’s independent Advisory Committee on Reactor Safeguards reaffirmed the importance of prototypes in new reactor development. Nevertheless, the NRC—a far weaker regulator today—has apparently changed its position and may proceed with licensing the ARDP demonstration reactors without requiring prototype testing first. But by skipping prototype testing and proceeding directly to commercial units, these projects may run not only the risk of experiencing unanticipated reliability problems, but also the risk of suffering serious accidents that could endanger public health and safety.

An additional challenge for NLWR demonstrations and subsequent commercial deployment is the availability of fuels for those reactors, which would differ significantly from the fuel that today’s LWRs use. Even a single small reactor could require a few tons of HALEU per year—far more than the 900 kilograms per year projected to be available over the next several years from a DOE-funded pilot enrichment plant that Centrus Energy Corporation is building in Piketon, Ohio. It is far from clear whether that pilot will succeed and can be scaled up in time to support the two NLWR demonstrations by 2027, not to mention the numerous other HALEU-fueled reactor projects that have been proposed.

The Future of the LWR

Those who argue that nuclear power’s progress depends on developing NLWRs have not made a persuasive case that the LWR has no future. LWR technology can realize nearly all the technological innovations attributed to NLWR designs, including passive safety features, the potential for modular construction, the use of advanced fuels, non-electric applications, greater plant autonomy to minimize labor costs, and underground siting. Although the LWR has its issues, NLWR designs clearly confront a different but no less formidable set of safety, security, and proliferation challenges.

A further consideration is how long it will take for new reactor types to achieve reliable performance once deployed. It took three decades for plant operators and researchers to increase the average capacity factor of the US fleet of LWRs from 50 to 90 percent. The relatively low state of maturity of NLWR technologies does not support the notion that these reactors will be able to achieve a similar level of performance in significantly less time.

Conclusions of the Assessment

The non-light-water nuclear reactor landscape is vast and complex, and it is beyond the scope of this report to survey the entire field in depth. Nevertheless, enough is clear even at this stage to draw some general conclusions regarding the safety and security of NLWRs and their prospects for rapid deployment.

Based on the available evidence, the NLWR designs currently under consideration (except possibly once-through, breed-and-burn reactors) do not offer obvious improvements over LWRs significant enough to justify their many risks. Regulators and other policymakers would be wise to look more closely at the nuclear power programs under way to make sure they prioritize safety and security. Future appropriations for NLWR technology research, development, and deployment should be guided by realistic assessments of the likely societal benefits that would result from the investment of billions of taxpayer dollars.

Little evidence supports claims that NLWRs will be significantly safer than today’s LWRs. While some NLWR designs offer some safety advantages, all have novel characteristics that could render them less safe.

All NLWR designs introduce new safety issues that will require substantial analysis and testing to fully understand and address—and it may not be possible to resolve them fully. To determine whether any NLWR concept will be significantly safer than LWRs, the reactor must achieve an advanced stage of technical maturity, undergo complete comprehensive safety testing and analysis, and acquire significant operating experience under realistic conditions.

The claim that any nuclear reactor system can “burn” or “consume” nuclear waste is a misleading oversimplification. Reactors can actually use only a fraction of spent nuclear fuel as new fuel, and separating that fraction increases the risks of nuclear proliferation and terrorism.

No nuclear reactor can use spent nuclear fuel directly as fresh fuel. Instead, spent fuel has to be “reprocessed”—chemically

treated to extract plutonium and other TRU elements, which must then be refabricated into new fuel. This introduces a grave danger: plutonium and other TRU elements can be used in nuclear weapons. Reprocessing and recycling render these materials vulnerable to diversion or theft and increases the risks of nuclear proliferation and terrorism—risks that are costly to address and that technical and institutional measures cannot fully mitigate. Any fuel cycle that requires reprocessing poses inherently greater proliferation and terrorism risks than the “once-through” cycle with direct disposal of spent fuel in a geologic repository.

Most NLWR designs under consideration do not offer obvious improvements over LWRs significant enough to justify their many risks.

Some NLWRs have the potential for greater sustainability than LWRs, but the improvements appear to be too small to justify their proliferation and safety risks.

Although some NLWR systems could use uranium more efficiently and generate smaller quantities of long-lived TRU isotopes in nuclear waste, for most designs these benefits could be achieved only by repeatedly reprocessing spent fuel to separate out these isotopes and recycle them in new fuel—and that presents unacceptable proliferation and security risks. In addition, reprocessing plants and other associated fuel cycle facilities are costly to build and operate, and they increase the environmental and safety impacts compared with the LWR once-through cycle. Moreover, the sustainability increases in practice would not be significant in a reasonably foreseeable time frame.

Once-through, breed-and-burn reactors have the potential to use uranium more efficiently without reprocessing, but many technical challenges remain.

One type of NLWR system that could in principle be more sustainable than the LWR without increasing proliferation and terrorism risks is the once-through, breed-and-burn reactor. Concepts such as TerraPower’s traveling-wave reactor could enable the use of depleted uranium waste stockpiles as fuel, which would increase the efficiency of uranium use. Although there is no economic motivation to develop more

uranium-efficient reactors at a time when uranium is cheap and abundant, reducing uranium mining may be beneficial for other reasons, and such reactors may be useful for the future. However, many technical challenges would have to be overcome to achieve breed-and-burn operation, including the development of very-high-burnup fuels. The fact that TerraPower suspended its project after more than a decade of development to pursue a more conventional and far less uranium-efficient SFR, the Natrium, suggests that these challenges have proven too great.

High-assay low enriched uranium (HALEU) fuel, which is needed for many NLWR designs, poses higher nuclear proliferation and nuclear terrorism risks than the lower-assay LEU used by the operating LWR fleet.

Many NLWR designs require uranium enriched to higher levels than the 5 percent U-235 typical of LWR fuel. Although uranium enriched to between 10 and 20 percent U-235 (defined here as HALEU) is considered impractical for direct use in nuclear weapons, it is more attractive for weapons use—and requires more stringent security—than the lower-assay enriched uranium in current LWRs.

The significant time and resources needed to safely commercialize any NLWR design should not be underestimated.

It will likely take decades and many billions of dollars to develop and commercially deploy any NLWR design, together with its associated fuel cycle facilities and other support activities. Such development programs would come with a significant risk of delay or failure and require long-term stewardship and funding commitments. And even if a commercially workable design were demonstrated, it would take many more years after that to deploy a large number of units and operate them safely and reliably.

Vendors that claim their NLWRs could be commercialized much more quickly typically assume that their designs will not require full-scale performance demonstrations and extensive safety testing, which could add well over a decade to the development timeline. However, current designs for sodium-cooled fast reactors and high-temperature gas-cooled reactors differ enough from past reactor demonstrations that they cannot afford to bypass additional full-scale prototype testing before licensing and commercial deployment. Molten salt-fueled reactors have only had small-scale demonstrations and thus are even less mature. NLWRs deployed commercially at premature stages of development run a high risk of poor performance and unexpected safety problems.

Recommendations

The DOE should suspend the advanced reactor demonstration program pending a finding by the NRC whether it will require full-scale prototype testing before licensing the two chosen designs as commercial power reactors.

The DOE has selected two NLWR designs, the Sodium SFR and the Xe-100 pebble-bed HTGR, for demonstration of full-scale commercial operation by 2027. However, the NRC has yet to evaluate whether these designs are mature enough that it can license them without first obtaining data from full-scale prototype plants to demonstrate novel safety features, validate computer codes, and qualify new types of fuel in representative environments. Without such an evaluation, the NRC will likely lack the information necessary to ensure safe, secure operation of these reactors. The DOE should suspend the Advanced Reactor Demonstration Program until the NRC—in consultation with the agency’s Advisory Committee on Reactor Safeguards and external experts—has determined whether prototypes will be needed first.

Congress should require that an independent, transparent, peer-review panel direct all DOE R&D on new nuclear concepts, including the construction of additional test or demonstration reactors.

Given the long time and high cost required to commercialize NLWR designs, the DOE should provide funding for NLWR R&D judiciously and only for reactor concepts that offer a strong possibility of significantly increasing safety and security—and do not increase proliferation risks. Moreover, unlike the process for selecting the two reactor designs for the Advanced Reactor Demonstration Program, decision-making should be transparent.⁶ Congress should require that the DOE convene an independent, public commission to thoroughly review the technical merits of all NLWR designs proposed for development and demonstration, including those already selected for the ARDP. The commission, whose members should represent a broad range of expertise and perspectives, would recommend funding only for designs that are highly likely to be commercialized successfully while achieving clearly greater safety and security than current-generation LWRs.

The DOE and other agencies should thoroughly assess the implications for proliferation and nuclear terrorism of the greatly expanded production, processing, and transport of the high-assay low-enriched uranium (HALEU) required to support the widespread deployment of NLWRs.

Large-scale deployment of NLWRs that use HALEU fuel will require establishing a new industrial infrastructure for producing and transporting the material. The DOE is actively promoting the development of HALEU-fueled reactor designs for export. Given that HALEU is a material of higher security concern than lower-assay LEU, Congress should require that the DOE immediately assess the proliferation and nuclear terrorism implications of transitioning to the widespread use of HALEU worldwide. This assessment should also address the resource requirements for the security and safeguards measures needed to ensure that such a transition can occur without an unacceptable increase in risk.

Congress should require that the DOE convene an independent, public commission to thoroughly review the technical merits of all NLWR designs proposed for development and demonstration.

The United States should make all new reactors and associated fuel facilities eligible for IAEA safeguards and provide that agency with the necessary resources for carrying out verification activities.

The IAEA, which is responsible for verifying that civilian nuclear facilities around the world are not being misused to produce materials for nuclear weapons, has limited or no experience in safeguarding many types of NLWRs and their associated fuel cycle facilities. NLWR projects being considered for deployment in the United States, such as the Sodium SFR and the Xe-100 pebble-bed HTGR, would provide ideal test beds for the IAEA to develop safeguards approaches. However, as a nuclear-weapon state, the United States is not obligated to give the IAEA access to its nuclear facilities. To set a good example and advance the cause of nonproliferation, the United States should immediately provide the IAEA with permission and funding to apply safeguards on all new US nuclear facilities, beginning at the design phase. This would help to identify safeguard challenges early and give the IAEA experience in verifying similar facilities if they are deployed in other countries.

The DOE and Congress should consider focusing nuclear energy R&D on improving the safety and security of LWRs, rather than on commercializing immature NLWR designs.

LWR technology benefits from a vast trove of information resulting from many decades of acquiring experimental data, analysis, and operating experience—far more than that available for any NLWR. This gives the LWR a significant advantage over other nuclear technologies. The DOE and Congress should do a more thorough evaluation of the benefits of focusing R&D funding on addressing the outstanding safety, security, and cost issues of LWRs rather than attempting to commercialize less mature reactor concepts. If the objective is to expand nuclear power to help deal with the climate crisis over the next few decades, improving LWRs could be a less risky bet.

EXECUTIVE SUMMARY ENDNOTES

- 1 *This report focuses on non-LWRs rather than LWR designs that differ from the operating fleet, such as the NuScale small modular reactor (SMR) design now under review by the Nuclear Regulatory Commission. UCS previously evaluated issues related to small modular LWRs in its 2013 report Small Isn't Always Beautiful. This report also does not discuss nuclear fusion reactors; despite some recent progress, these likely remain even further away from commercialization than the early-stage fission reactor concepts.*
- 2 *Some sources define HALEU as LEU enriched from 5 percent to less than 20 percent uranium-235. However, this range does not align with the nuclear security risk of different grades of LEU. This report adopts the definition of HALEU used by the uranium enrichment consortium URENCO.*
- 3 *In any event, non-LWRs do not have a monopoly on non-electric applications. Current-generation LWRs as well as small modular LWRs are being piloted for non-electricity applications such as producing hydrogen. At least one type of novel LWR, the super-critical LWR, would be capable of producing high-temperature steam, but it is not currently under development.*
- 4 *One way to do that would be for the United States to designate all new reactors and fuel cycle facilities as eligible for IAEA safeguards. This would give the agency an opportunity to develop verification approaches for new facility types—if such approaches are feasible.*
- 5 *The treaty names 44 countries that must sign and ratify it before it enters into force. To date, eight of these countries have not ratified and/or signed the treaty—including the United States, which has signed but not ratified it.*
- 6 *Although the DOE has said that an external review of its selections took place, it has not publicly released the reviewers' names and affiliations—nor has it publicly documented their findings.*

Nuclear Power: Present and Future

The future of nuclear power is uncertain, both in the United States and worldwide. Because nuclear power is a low-carbon way to generate electricity, there is considerable interest in expanding its role to help mitigate the threat of climate change. However, the technology has fundamental safety and security disadvantages compared with other low-carbon sources. The nuclear industry, policymakers, and regulators must address these shortcomings fully if the global use of nuclear power is to increase around the world without posing unacceptable risks to public health, the environment, and international peace and security.

Almost all nuclear power reactors operating today are light-water reactors (LWRs), so called because they use ordinary water (H₂O) as a coolant.¹ Of the approximately 50 power reactors under construction around the world, all but a few are water-cooled. Most new projects are conventionally sized large reactors with power production capacities of at least 3400 megawatts of thermal energy (MWth), equivalent to about 1,100 megawatts of electricity (MWe).

Slower Growth, Cost and Safety Concerns

Despite renewed enthusiasm for nuclear power in many quarters, its recent growth has been far slower than anticipated ten years ago amid continuing debate over its risks, costs, and benefits (Schneider et al. 2020). At the end of 2010, there were 441 operating nuclear power reactors worldwide, with a total electrical power capacity of 375 gigawatts of electricity (GWe) (IAEA 2011). At the end of 2019, there were 443 operating reactors—only two more than in 2010—with a total generating capacity of 392 GWe (IAEA 2020). In 2019, these reactors generated 2,657 terawatt-hours of electricity,

or 10.4 percent of total electricity generation. This actually represented a decrease of over 20 percent in the share of the global electricity demand met by nuclear energy compared to 2010 (IAEA 2020).

There are a number of reasons why global nuclear power capacity and its share of electricity demand has not increased over the last decade, despite prior expectations of a so-called nuclear renaissance. The Fukushima Daiichi triple nuclear reactor meltdown in Japan in 2011 no doubt played a role in slowing down nuclear power expansion in some countries. The accident contaminated a wide area with long-lived radioactivity and led to a prolonged shutdown of Japan's other nuclear plants. Japan's nuclear sector may not return to its level prior to the accident for decades, if ever. The accident caused some countries, such as China, to temporarily pause nuclear plant construction, and even prompted some others, such as Switzerland, to decide to phase out nuclear power entirely.

However, the Fukushima accident has not proven to be a decisive consideration for most other countries' energy programs. A more significant factor affecting nuclear energy's prospects is its high cost today relative to other sources of electricity—primarily natural gas but also, increasingly, renewable energy sources such as wind and solar.

NEW NUCLEAR: PROHIBITIVELY EXPENSIVE

In the United States, new nuclear plant projects have proved to be prohibitively expensive and have lengthy construction times, discouraging private investment. The only reactors being built in the United States today, two 1100 MWe Westinghouse AP1000 units at the Vogtle plant in Georgia, are now projected to cost nearly \$14 billion each and will take

at least a decade to complete—twice the cost and more than twice as long as original estimates (Nuclear Engineering International 2020). (This estimate predated the 2020 coronavirus pandemic, which has caused further delays.) Similar problems have plagued the 1600 MWe reactors of Areva’s EPR design being built at Olkiluoto in Finland, Flamanville in France, and Hinkley Point in the United Kingdom. In the United States, even some nuclear plants already operating are costlier to run than natural gas plants and new wind and solar projects, and are being retired before they reach the end of their service lives (Clemmer et al. 2018).

Most reactors under construction today are in countries that provide substantial government support, such as China, and are less susceptible to market pressures. Even for such countries, however, cost is still a factor. Government treasuries are not unlimited, and nuclear power subsidies must compete with other uses of public funds. For example, China’s nuclear power growth has slowed, and its total nuclear capacity at the end of 2020, around 50 GWe, fell short of its aggressive target of 58 GWe.

Assuming that market conditions do not change significantly, the future of nuclear energy around the world over the next few decades will depend in large part on national decisions about the role nuclear power should play in addressing climate change and the extent to which governments should underwrite them. In its 2020 annual report, the International Atomic Energy Agency (IAEA) estimated that by 2050, the change in world nuclear-energy generating capacity could range from an increase of 82 percent for its “high case” projection to a decrease of seven percent in the “low case” (IAEA 2020). The IAEA’s high case corresponds to an average annual increase of around 10 GWe per year (about 10 conventionally sized reactors)—lower than past projections but still higher than recent annual growth.² However, since 2014, an average of fewer than five new reactor construction projects have started up per year (IAEA 2020).

Most of the IAEA’s projected growth in nuclear power capacity is in developing countries, with stagnant or declining capacity in industrialized nations. In the United States, the Energy Information Administration’s Annual Energy Outlook 2021 projects nuclear power capacity to decline from the 2019 level of 98 GWe to about 75 GWe by 2050 under a reference case with no change in current energy policies, and to about 42 GWe by 2050 under a low natural gas-price case (EIA 2021).

Therefore, without national policies to limit carbon emissions, nuclear power will likely remain at a competitive disadvantage as long as fossil fuel prices remain low. Mechanisms for internalizing the social cost of carbon, such as a carbon tax, could help level the playing field for low-carbon energy sources. But it is not clear that nuclear power would thrive

even with a high price on carbon, which would also benefit other low-carbon electricity sources that have fewer safety and security problems. Other government actions likely would be necessary for nuclear power to expand enough to make a significant contribution to reducing carbon emissions in the next few decades and help mitigate the most severe impacts of climate change. These actions include substantial, long-term investments in the supporting infrastructure, strong safety requirements, and credible plans for disposal of long-lived nuclear wastes.

However, some observers believe that the technology itself is the problem and that more radical fixes are needed (Soltoff 2020). They argue that the large LWR, the nuclear industry’s workhorse, has inherent flaws that inhibit nuclear power’s growth. In addition to these reactors’ high costs and long construction times, they point to (among other things) LWRs’ susceptibility to severe core-melt accidents such as occurred at Fukushima, their inefficient use of uranium, and the long-lived nuclear wastes they generate.

No matter what the root causes, the poor image projected by troubled LWR construction projects such as Vogtle and Olkiluoto has no doubt contributed to a credibility problem for an industry that promotes nuclear power as the world’s best hope for mitigating climate change based on assertions that it is affordable and can be quickly deployed on a large scale (WNA, n.d.). One response to public skepticism about the current state of nuclear power is for developers to pursue different types of reactors—some radically different—that they promise will be safer, cheaper, and quicker to build. But a fundamental question about these alternative designs remains: Is *different* actually *better*? This report aims to shed light on that question.

Can Non-Light-Water Reactors Revive Nuclear Power’s Prospects?

There are three main strategies for shifting the current nuclear power paradigm, with the aim of mitigating the technology’s safety, sustainability, and cost problems.

The first approach is to develop new types of large LWRs that would be safer while also being cheaper to build and operate. This is the path that Westinghouse and Areva respectively pursued—arguably unsuccessfully—with their AP1000 and EPR designs. In parallel, new types of “accident tolerant” fuels for LWRs that in principle could reduce the risk of melt-down are being developed in several countries, including the United States. However, data remains sparse, and early results have not been promising (Khatib-Rahbar et al. 2020).

The second approach is to go small. Some observers believe the future lies in small, “modular” reactors with

capacities of 300 MWe or less, known as SMRs. Small modular LWRs could be somewhat safer than large LWRs by virtue of their size and lower rate of heat production, but they would produce more expensive electricity without employing measures to significantly cut capital and operating costs per megawatt (Lyman 2013). These reactor modules could be mass-produced in a factory and deployed as needed to meet electricity demand growth, either singly or in groups—features that proponents argue could lower construction and financing costs. A number of small modular LWRs, such as the 77 MWe NuScale reactor, are currently in development. While these designs have some novel features, they are essentially evolutionary variants of current LWRs.

The third is to go in a new direction and develop reactors that are not cooled by water but by other substances, such as liquid sodium, helium gas, or even molten salts. Such reactors, known as non-light-water reactors (NLWRs), would differ from LWRs in many fundamental aspects. Numerous vendors, both established companies and small startups, are pursuing development of NLWR technologies.³

In general, either LWRs or NLWRs can be SMRs. Proposed NLWRs range from units as large as today’s operating reactors to “micro-reactors” with capacities of less than 10 MWe. Some NLWR designs have capacities of 300 MWe or less and therefore qualify as SMRs, but others do not, as they must be a certain minimum size to work effectively.

NLWR developers state variously that their designs have the potential to lower cost, reduce the accumulation of nuclear waste, use uranium more efficiently, improve safety, and reduce the risk of nuclear proliferation (see, for example, Back 2017). More specifically, they cite features such as modular construction, passive safety, underground siting, and—for some designs—the ability to provide high-temperature process heat for manufacturing. Some vendors promise that their designs can be demonstrated, licensed, and deployed on a large scale within a decade or two.

Are these claims justified? How can one identify genuine innovations amongst the hype? As with any new technology, an independent reality check is needed. From self-driving cars to finger-prick blood tests to cheap flights to Mars, the Silicon Valley-style disruptive digital technology model has not always proven readily adaptable to other engineering disciplines. And nuclear energy, which requires painstaking, time-consuming, and resource-intensive research and development, is proving to be one of the harder technologies to disrupt.

A Note on Terminology

Reactor concepts that differ from conventional LWRs are often referred to as “advanced” reactors, although this definition is not universally used from one government agency to another or even from one piece of federal legislation to another. Most recently, Congress defined an advanced reactor as “any light water or non-light-water fission reactor with significant improvements compared to the current generation of operational reactors” (Energy and Water Development and Related Agencies Appropriations Act of 2020, Pub. L. No. 116-94, 133 Stat. 2535 (2019)).

To avoid confusion, this report will not use the term “advanced reactor.” The present report focuses on non-light-water reactors (NLWRs) (see endnote 1).

The Need to Fully Vet Claims About NLWRs

If nuclear power is to succeed in the future, it is critical that government policymakers and private investors fully vet the claims new reactor developers are making and accurately assess their prospects for successful development and safe, secure, and cost-effective deployment. Given the urgency of the climate crisis, these technologies need to be rigorously evaluated to avoid wasting time and resources on concepts that are high-risk but would offer only low potential benefits in practice. Weeding out such technologies would help researchers focus on other approaches to climate mitigation that are less risky and more beneficial.

Key questions that policymakers should consider are the following:

- Do NLWRs offer significant benefits over LWRs?
- How do the safety, proliferation, and environmental risks of NLWRs compare to those of LWRs?
- Do the potential benefits of NLWRs outweigh the risks and justify the substantial public and private investment needed to commercialize them?
- Can NLWRs be safely and securely commercialized in time to contribute significantly to averting the climate crisis?

The purpose of this report is to help inform policy decisions on these questions by critically evaluating certain claims being made about each of the principal classes of NLWRs: liquid metal-cooled fast reactors, high-temperature gas-cooled reactors, and molten salt-fueled reactors.⁴ In particular, it compares NLWRs to LWRs in terms of safety, security, nuclear proliferation risk, and sustainability—the latter including the often-claimed ability of these reactors to

“recycle” nuclear waste to reduce the amount requiring long-term geological isolation. The report also considers “breed-and-burn” reactors, which, in theory, could use uranium fuel much more efficiently without the need to recycle their spent fuel.

Each NLWR design has both advantages and disadvantages compared to LWRs. New nuclear reactor types that offer significant safety, security, or economic benefits compared to LWRs would be welcome—as long as improvements in one area did not cause greater problems in others. But this is where the design challenges lie. For example, increasing the size of a reactor core while keeping its power capacity fixed could make the core easier to cool in an accident but might also increase cost.

There is certainly room for innovation in nuclear technology. For example, advances in materials science can increase the durability of reactor structures and fuels. Faster computation can improve the modeling of reactor operation. More efficient cooling system designs can reduce or remove the need for large volumes of water to generate steam—an important consideration in light of increasing surface water temperatures and increasing water scarcity resulting from climate change. And changing objectives—for instance, the need to prevent hydrogen explosions such as those that destroyed three reactor buildings at Fukushima—can stimulate new approaches to solving old problems.

However, a quicker payoff is more likely to be achieved by focusing research on improving well-established reactor technologies than by pursuing the development of speculative designs that have hit roadblocks in the past and have had little or no operating experience as a result. It took three decades for plant operators and researchers to increase the average capacity factor of the US fleet of LWRs from 50 percent to 90 percent by correcting problems that affect reliability, such as coolant-material interactions. The relatively low state of maturity of NLWR technologies does not support the notion that these reactors will be able to achieve a similar level of performance in significantly less time.

NLWRs: Past and Present

Another reason to avoid the term “advanced reactor” is that it is a misnomer for most of the designs being pursued today, which are based on decades-old concepts. As an Oak Ridge National Laboratory scientist put it succinctly in a January 2019 presentation, “today’s ‘advanced reactors’ closely resemble their 1950s–1970s predecessors in: core configuration; materials in structure, core, and fuel; approach to [fuel] qualification; and control systems” (Terrani 2019).

Much of the creativity in nuclear plant design dates back to the 1940s, the early years of the nuclear power era, when Manhattan Project scientists and engineers engaged in wide-ranging brainstorming to explore the full potential of the new nuclear technology. One “advanced” reactor design, the liquid metal-cooled fast-neutron reactor, even predates the forerunner of today’s LWR. For decades in the mid-20th century, the Atomic Energy Commission (the predecessor of today’s Department of Energy and Nuclear Regulatory Commission) devoted considerable time and resources to developing a variety of LWR and NLWR technologies, demonstrating many designs at various scales at sites around the United States. A number of other countries also built and operated NLWRs (see Table 1, p. 17).

Most of the prototype reactors encountered operational problems, and some even experienced serious accidents. To be sure, LWRs experienced accidents as well—most notably the Three Mile Island meltdown in 1979. However, over time the LWR became predominant. To some extent, this was because the LWR, the design chosen originally by the Navy in the 1950s for submarine propulsion, received much more funding than other designs. But several of the LWR’s rivals were abandoned after experiencing engineering challenges that proved too difficult to overcome.

Utilities that gambled on NLWRs ultimately lost their bets. A consortium led by Detroit Edison built a small sodium-cooled fast reactor in Michigan, called Fermi-1. Soon after reaching full power, the reactor partially melted down in 1966, and did not restart until 1970, only to be shut down for good in 1972. The Fort St. Vrain reactor, a high-temperature gas-cooled reactor built by the Public Service Company of Colorado, operated for only a decade, with an average capacity factor of only 14 percent, before being shut in 1989. It is reasonable to surmise that LWRs emerged as the industry standard because they simply proved better suited to meet the needs of utilities and consumers.

CURRENT DEPARTMENT OF ENERGY WORK ON NLWRs

Given the problems with current-generation LWRs, the development of new nuclear technologies that can significantly increase safety and security while being more cost-effective is a worthwhile goal. However, it is not clear that current nuclear energy programs are being designed to make this objective a priority.

The US government has continued to conduct research and development on various types of NLWRs and their fuel cycles.⁵ When it was created in 1977, the Department of Energy (DOE) inherited the former Atomic Energy Commission’s nuclear power portfolio, and the DOE has continued

TABLE 1. Past and Present Demonstration Reactors Worldwide

Step in Deployment Path	Light-Water Reactor (example)	Sodium-Cooled Fast Reactor		High-Temperature Gas-Cooled Reactor		Molten Salt-Fueled Reactor	
	US	US	Intl	US	Intl	US	Intl
R&D for scientific feasibility	SPERT BORAX PBF	SEFOR (20 MWth) TREAT	CABRI ^a	None	None	None	None
Engineering Demonstration	S1W EBWR	EBR-I (1.4 MWth) EBR-II (20 MWe)	Dounreay (14 MWe) Rhapsodie (40 MWth)	Peach Bottom (40 MWe)	DRAGON (20 MWth) HTR-10 ^a (10 MWth) HTTR ^a (30 MWth) AVR (15 MWe)	Aircraft Reactor Experiment (2.5 MWth) MSRE (7.4 MWth)	None
Performance Demonstration	USS Nautilus Shippingport	Fermi-1 (69 MWe) FFTF (400 MWth)	CEFR ^a (65 MWth) Phénix (233 MWe) Monju (300 MWe) BN-300 (300 MWe) BN-600 ^a (600 MWe) PFR (250 MWe)	FSV (842 MWth) ^b	THTR (750 MWth) ^b	None	None
Commercial Demonstration	Yankee Rowe (485–600 MWth)	None	Superphénix (3000 MWth) BN-800 ^a (800 MWe)	None	None	None	None

a Reactor is still operational as of February 2021.

b FSV and THTR were commercial demonstrations of large HTGRs; however, for modular HTGRs under consideration today, they serve the role of a performance demonstration.

SOURCE: ADAPTED FROM PETTI ET AL. 2017.

to pursue multiple NLWR designs. In the 1990s, it initiated the Generation IV program, with the goal of “developing and demonstrating advanced nuclear energy systems that meet future needs for safe, sustainable, environmentally responsible, economical, proliferation-resistant, and physically secure energy” (INL 2005).

Under Generation IV, the DOE identified six families of reactor technologies including five NLWRs and one LWR: the sodium-cooled fast reactor, lead-cooled fast reactor, molten salt reactor, gas-cooled fast reactor, supercritical LWR, and very high temperature gas-cooled reactor. The DOE’s

funding priorities have varied over the years depending on congressional mandates and internal competition, but most support has gone to the development of sodium-cooled fast reactors and high-temperature gas-cooled reactors. Around 2008 the DOE resumed funding molten salt reactor development, after a hiatus of several decades.

More recently, with strong support from Congress, the DOE has expanded its NLWR activities and is pursuing several new reactor projects. It is proceeding with the design and construction of a sodium-cooled fast reactor called the Versatile Test Reactor (VTR), which it hopes to begin

operating by 2026–2031. The VTR would not generate electrical power but would be used to test nuclear fuels and materials for development of other reactors. The DOE has also selected designs for two NLWRs—a high-temperature gas-cooled reactor and a second sodium-cooled fast reactor—with the intent of demonstrating them for commercial power production by 2027. And the DOE and the Department of Defense are both pursuing demonstrations of so-called micro-reactors—that is, NLWRs with capacities from 1 to 20 MWe—also within the next few years. However, Congress will need to provide sufficient and sustained funding for any of these projects to come to fruition. Appropriated funding for the VTR to date has already fallen far short of the amount that the DOE has requested to support its current schedule.

THE OBJECTIVES OF ADVANCED REACTOR DEVELOPMENT

In the FY 2020 Energy and Water Development Appropriations Act, Congress defined an advanced reactor as

any light water or non-light-water fission reactor with significant improvements compared to the current generation of operational reactors. Significant improvements may include inherent safety features, lower waste yields, greater fuel utilization, superior reliability, resistance to proliferation, increased thermal efficiency, and the ability to integrate into electric and nonelectric applications (Energy and Water Development and Related Agencies Appropriations Act of 2020, Pub. L. No. 116-94, 133 Stat. 2535 (2019)).

Some of these goals—namely, increasing safety and reliability—are worthwhile. Others, such as greater fuel utilization, are less compelling, given the abundance of natural uranium for fuel (discussed below). One problem with the above statutory definition of an advanced reactor is that new reactor designs may have significant improvements in one or more categories but also significant disadvantages in others—drawbacks that may outweigh the benefits. It is important that policymakers consider the full spectrum of positive and negative attributes of new designs to identify reactors with the highest potential for significant improvements *overall* compared to current technologies. Otherwise, the deployment of a new type of reactor may create more problems than it solves.

Unfortunately, in its 2017 assessment of NLWRs, the DOE did not consider the full range of objectives for developing these technologies, but chose to focus on only two. The first is to “deploy a high-temperature process heat application” for industrial activities such as synfuels production, and the second is to “extend natural resource utilization and

reduce the burden of nuclear waste for future generations” (Petti et al. 2017). Notably, the report did not stress the importance of other considerations, including safety, security, proliferation resistance, or economics.

How important are the two performance objectives that DOE considered in its assessment for advancing nuclear power?

PROCESS HEAT

The goal of developing nuclear reactors to provide industrial process heat does not appear to be driven by demand from the industrial sector. Although the nuclear industry has been pushing the idea of developing high-coolant-temperature reactors for non-nuclear process heat applications for decades, there is little evidence that the industries that would utilize such heat are themselves interested in using nuclear power. And it is unclear why these other industries would want to incur the additional risks of operating nuclear reactors in proximity to chemical plants.

A 2004 report by the Nuclear Energy Agency stated that “the reality does not match the potential,” and posed the question “if nuclear energy has so high potential in the non-electricity product market, why has its deployment been so limited? Can one expect some dramatic changes in this market situation?” (NEA 2004). Virtually the same question was asked nine years later, in a 2013 joint NEA/IAEA workshop (Paillère 2013). And in 2018, the IAEA reported that experts at a meeting on the subject agreed that “for these . . . products to enter the commercial market on a large scale, several challenges and barriers have to be overcome,” including economics, low public acceptance, and technical and regulatory issues (Dyck 2018). Apparently there remains little interest in these applications by potential users.

Nevertheless, industrial processes are a significant contributor to carbon emissions, and the economics of nuclear process heat would improve with a price on carbon.⁶

SUSTAINABILITY

The importance of the second objective cited in the DOE’s 2017 assessment, often referred to as sustainability, is also questionable.

There are two primary aspects to improving sustainability relative to current-generation LWRs. The first is increasing the efficiency of use of natural resources (e.g., mined uranium), and the second is reducing the quantity of long-lived, heat-generating radionuclides contained in radioactive waste and that need to be disposed of in a geologic repository (primarily plutonium and other transuranic elements).

Some advanced reactor developers have taken this concern about sustainability to an extreme, invoking mislead-

ing—but very compelling—messages about the ability of their designs to “consume,” “burn,” or “recycle” spent fuel from LWRs. For example, Jacob DeWitte, co-founder of Oklo, Inc., testified before Congress that his fast reactor concept “can consume the used fuel from today’s reactors” (DeWitte 2016). According to the GE-Hitachi website, the PRISM (Power Reactor Innovative Small Module) fast reactor and its associated reprocessing facility would “recycle all the uranium and transuranics . . . contained within used nuclear fuel” (GEH 2021). General Atomics states that its high-temperature gas-cooled fast reactor is capable of turning our waste stockpile into an important energy resource (GA 2019). And until shortly before it shut down in 2018, the company Transatomic Power claimed, erroneously, that waste from conventional nuclear reactors could be used as the fuel for its MSR.

The reality is much more complicated. First, these statements greatly exaggerate the actual capabilities of these reactors to achieve these goals. Second, for any reactor concept it is critical to understand that “burning” spent fuel first entails reprocessing to separate out and re-use plutonium and other weapon-usable materials. Reprocessing makes these materials more accessible for use in nuclear weapons by states or terrorists, as explained below.

In theory, some NLWRs could make more efficient use of uranium or waste repository capacity. Indeed, greater sustainability is one of the only clear advantages that certain NLWRs, such as fast reactors, could offer over LWRs. However, for such reactors and their fuel cycles, it has not been established that the real-world benefits would be large enough to justify their proliferation and safety risks, not to mention their enormous development costs.

The possible exceptions are once-through “breed-and-burn” reactors, which have the potential to use uranium more efficiently than LWRs without reprocessing and recycling spent fuel. If these reactors could be successfully developed, they would remove one of the major incentives cited by advocates for reprocessing—and avoid the associated risks. However, these concepts are proving difficult to realize, and they have safety problems and other challenges. It is not clear whether such designs will be feasible.

REDUCING NUCLEAR PROLIFERATION AND NUCLEAR TERRORISM RISKS: A FUNDAMENTAL OBJECTIVE

The failure to consider other important criteria, such as safety and nuclear proliferation, was a significant shortcoming of the DOE’s 2017 study. In particular, the assessment was skewed by its lack of attention to proliferation and terrorism risks, which are critical issues for evaluating alternative reactors and fuel cycles. Fuel cycles that involve reprocessing and recycle of nuclear-weapon-usable materials may offer

increased sustainability compared to LWRs operating on a once-through cycle without reprocessing, but they will pose greater security risks —risks that design and operational features cannot fully mitigate.

Therefore, when assessing the overall benefits of NLWRs and their fuel cycles, it is important to demonstrate meaningful overall benefits in safety, security, and cost-effectiveness, and that do not involve reprocessing.

A Host of Challenges Even for More Mature NLWR Designs

How much time would it take to commercialize a novel NLWR concept? Some NLWR developers say that they will be able to deploy their reactors commercially as soon as the late 2020s. However, such timelines are not likely to be realistic, and could only be met by bypassing many of the developmental stages necessary for ensuring safe and secure operation.

The DOE has identified four stages to fully develop a reactor design that has not been built before (Box 1) (Petti et al. 2017). The pathway includes construction and operation of one reduced-scale and two full-scale prototype reactors

BOX 1.

Stages of Advanced Reactor Development

- **Research and development to prove scientific feasibility** of key features associated with fuel, coolant, and geometrical configuration. Irradiation test reactor services are particularly important in this phase, although they can be beneficial at each step (e.g., to explore additional fuel/material options).
- **Engineering demonstration at reduced scale** for proof of concept for designs that have never been built. The goal at this demonstration level is to test the viability of the integrated system. Historically, these have been small reactors (less than 50 MWe).
- **Performance demonstration(s) to establish that scale-up** of the system works and to gain operating experience to validate the integral behavior of the system (including the fuel cycle in some cases), resulting in proof of performance.
- **Commercial demonstrations** that will be replicated for subsequent commercial offerings if the system works as designed.

SOURCE: PETTI ET AL. 2017

(which the DOE refers to as “demonstration” reactors). Carrying out this program in its entirety could take several decades and cost many billions of dollars.

Prototype testing is needed, among other things, to confirm that reactor systems will work as intended, to demonstrate reliability of the reactor as a whole, to qualify reactor fuels, and to assess the effectiveness of new safety features. In the United States, the Nuclear Regulatory Commission (NRC) generally requires prototype testing for licensing designs that differ significantly from LWRs, unless the agency determines that such testing is not needed because sufficient data exists from test programs, analyses, and past experience. If a reactor applicant chooses to test a prototype, the NRC may impose additional safety requirements to protect public health and safety that would not apply to a commercial reactor.

A key consideration in estimating the commercialization timeline for designs that have had some previous research and development is their technical maturity—in particular, whether sufficient historical data exists to validate design features and analysis methods and enable developers to leapfrog over one or more stages, such as engineering or performance demonstrations (see Box 1, p. 18).

In a 2017 study, the DOE estimated the time and resources needed to commercialize different types of NLWRs (Petti et al. 2017). The DOE judged that two NLWR categories—sodium-cooled fast reactors and high-temperature gas-cooled reactors—were sufficiently mature that they did not require additional engineering or performance demonstrations before proceeding to commercial demonstration. Based on vendor-supplied data, the DOE estimated that it would cost approximately \$4 billion and take 13 to 15 years to begin operating the first commercial demonstration unit of either type, assuming that many aspects of the project could proceed in parallel, such as technology development, design, and licensing. Reactor construction and startup testing was assumed to take seven years of the 13- to 15-year period. After five years of operation of the commercial demonstration unit, additional commercial deployments could follow in the “2030 time-frame” (Petti et al. 2017).

In contrast, for lower-maturity designs such as molten salt-cooled reactors, the DOE report concluded that both engineering and performance demonstration reactors would be needed—stages that could cost an additional \$2 billion to \$4 billion and add 20 years to the timeline. The subsequent commercial demonstration reactor would also cost billions of dollars and would not begin until 2040, and the model would not be available for sale until the mid-2040s or even the 2050s.

Although these timelines are long and the costs are high, they are likely too optimistic, and inconsistent with the recent experience of new reactors such as the Westinghouse AP1000, an evolutionary LWR. Although the AP1000 has some novel features, it is fundamentally based on mature LWR technology. Therefore, Westinghouse was able to leverage many decades of LWR operating data, and the company did not build a prototype reactor before licensing and selling commercial units. Even so, it took more than 30 years of research, development, and construction before the world’s first AP1000 unit—Sanmen-1 in China—began to produce power in 2018. The first US AP1000 unit, the Vogtle-3 reactor in Georgia, is taking even longer, and is not slated to begin operation before November 2021.

Congressional supporters of new nuclear reactor development are determined to speed up the process. In 2020, Congress created the Advanced Reactor Demonstration Program (ARDP) within the DOE to accelerate commercialization of new reactor types at different stages of maturity. The ARDP has provided initial cost-shared awards to industry to build two commercial demonstration “advanced” reactors in only five to seven years, or by 2027 at the latest—far less time than the DOE’s earlier estimate of at least 13 to 15 years for the more mature designs. The ARDP currently caps the total DOE contribution at \$1.6 billion per design for both the reactors and supporting fuel production facilities (for a total of up \$3.2 billion each, with the vendors contributing 50 percent). The ARDP is also providing an additional funding stream for development of less mature designs, such as molten salt reactors, with the expectation they will have an operational reactor within 10 to 12 years—again, far short of the 25–30 years that the DOE previously estimated. (All future ARDP funding is subject to congressional appropriations.)

In October 2020, the DOE, consistent with its 2017 assessment of which NLWR types are most mature, chose the Natrium, a 345 MWe (or 840 MWth) sodium-cooled fast reactor being developed by TerraPower, and the (approximately) 76 MWe Xe-100, a high-temperature gas-cooled reactor being developed by X-Energy, for the ARDP commercial demonstrations. (The Xe-100 will be deployed in a four-pack for a total of about 300 MWe.) Because these reactors will generate commercial power, the Atomic Energy Act requires that they be licensed by the NRC.

In parallel with these projects, the DOE is also planning to build the Versatile Test Reactor (VTR), a 300 MWth sodium-cooled fast reactor based on the same fundamental design as the Natrium. Unlike the Natrium, however, the VTR would not generate electrical power but would be used to test

materials and fuels for other fast reactor designs—and would not require NRC licensing.

Even the DOE admits that the ARDP timeline—which falls short of its earlier, 13- to 15-year projection for commercialization of mature NLWR designs—is very ambitious. And the agency’s experience to date with the VTR should sound a cautionary note regarding schedule and cost predictions for NLWR development. In 2019, the DOE projected that the VTR could be built and started up by 2025 and cost up to \$4.5 billion. However, the agency soon admitted that construction may take until 2031 and could cost up to \$5.8 billion. Even this is probably an underestimate, given the DOE’s poor track record for making large capital project schedule and cost projections.

Moreover, it is not clear that even the longer, 13- to 15-year development timeline is realistic for the Sodium and Xe-100 designs that the DOE chose for the ARDP. As discussed in chapters 5 and 6, past sodium-cooled fast reactor and high-temperature gas-cooled reactor demonstrations have had safety and reliability problems. In addition, both of these designs differ significantly from those earlier demonstration reactors in ways that are important to safety. The safety of these commercial demonstration reactors could well be in question if they are built and operated without prior prototype testing under controlled conditions.

In the 1990s the NRC concluded, after reviewing available data from prior demonstrations, that it would require representative prototype testing before licensing either a sodium-cooled fast reactor or a high-temperature gas-cooled reactor. However, the NRC is a less safety-focused agency today and may relax its requirements. Nevertheless, to license the two ARDP reactors, the NRC will need to soon determine whether prototype testing will be necessary—a decision that could significantly affect project costs and schedules. The NRC has encouraged NLWR applicants to develop regulatory plans prior to licensing to engage the agency as early as possible regarding their intentions regarding prototype testing (NRC 2017), but there is no indication that the prospective applicants have submitted such plans yet to the NRC. Thus, there is considerable uncertainty whether the ARDP goal of commercial reactor operation by 2027 is compatible with the NRC’s obligation to ensure protection of public health and safety.

COSTLY AND LENGTHY DEVELOPMENT OF FUEL INFRASTRUCTURE

In addition to the development of reactor technology, the commercialization of a new nuclear plant design requires development of the associated fuel cycle infrastructure,

which would also be a costly and lengthy undertaking. New facilities will be needed to fabricate novel types of nuclear fuel for the reactors and to manage their spent fuel, and the current system for nuclear fuel and waste transportation would need to be modified to handle materials with different characteristics.

Indeed, perhaps the biggest challenge for near-term NLWR demonstrations and subsequent commercial deployment is the availability of fuels for those reactors, which would be significantly different from the fuel used by LWRs today. In particular, many proposed reactors, including both the Sodium and Xe-100 demonstration reactors, would need large quantities of uranium enriched to higher levels—so-called high-assay low-enriched uranium (HALEU). HALEU is a material that is in very short supply and not commercially available. Even a single small reactor could require tons of HALEU per year (see chapter 4), far more than the current available supply. The Nuclear Energy Institute has estimated that it would take a minimum of seven to nine years to establish a domestic fuel-cycle infrastructure to support a significant level of HALEU production, assuming full funding is available (NEI 2018).

However, this funding assumption is questionable. The only operating US enrichment plant, URENCO-USA in New Mexico, has expressed willingness to produce HALEU but has not made any commitment to proceed in the absence of a strong market signal that demand will materialize. The company has called for “sustained and dedicated” government funding for such a program and has proposed that the DOE become a wholesale buyer of HALEU, at least for the initial output (Fletcher 2020). However, the only near-term prospect for production of HALEU is a three-year pilot centrifuge enrichment demonstration project the DOE has sponsored at the Centrus Energy Corporation facility in Piketon, Ohio, but that will produce, at most, a few hundred kilograms by June 2022. Centrus estimates that the facility could eventually produce up to around 900 kilograms of HALEU per year—not nearly enough for the demonstration reactors (Dyke 2020).⁷

As new types of fuels are developed and produced, they must undergo rigorous qualification programs before they can be safely used in reactors—also a time-consuming and costly process. The former director of the DOE’s high-temperature gas-cooled reactor fuel development program, Dr. David Petti, has been candid about the considerable time and resources needed to fully qualify new types of fuel, which is a painstaking and slow process that can involve trial and error. The process, in which fuel samples are irradiated under representative conditions, cooled, analyzed in detail, and subject to transient testing to simulate accidents, may need to be

repeated several times. Once a fuel form has been established, it would take between 15 and 25 years to complete the development process, with the lower bound unrealistically assuming a program with no resource constraints and in which everything proceeds as expected (Petti 2016).

SUSTAINED GOVERNMENT FUNDING REQUIRED

The various cost and time projections for commercializing NLWRs may differ in the details, but they all illustrate the significant technical challenges encountered in developing a new reactor design and its associated fuel cycle. Even TerraPower—likely the best funded reactor startup—was apparently unwilling to spend the many billions of dollars needed to commercialize its concepts on its own, and did not move forward with a demonstration reactor until it had secured government funding through the ARDP.

Thus, to commercialize any NLWR design, the ARDP example shows that government will likely need to provide substantial and sustained funding—not only for fundamental research, development, and demonstration, but perhaps even for the deployment of the first commercial units. As a 2014 DOE study concluded, “the market disincentives and barriers to commercial implementation of nearly all the promising [NLWR] options are expected to be very significant, such that federal government intervention . . . will likely be required for full-scale implementation of a new fuel cycle . . .” (Wigeland et al. 2014).

Although new nuclear technologies may not be attractive to investors looking for short-term returns, they may have longer-term societal benefits. For instance, NLWRs that are more costly than LWRs but use uranium more efficiently might help ensure future resource availability. However, they would not be a good choice for a utility as long as there is a cheap and plentiful fuel supply—as is the case now and for the foreseeable future. Similarly, utilities would have no incentives to choose a safer reactor that would cost more and exceed regulatory requirements.

Thus, government support for NLWR development could be justifiable—but only for designs with a high likelihood of significantly advancing nuclear power technology in multiple areas. Developing reactors and fuel cycles that would only offer marginal improvements over LWRs, or that would increase safety, security, or proliferation risks, are not wise uses of taxpayer funds.

Nuclear Power Growth and Climate Change Mitigation

The timeline for commercialization of NLWRs is a key factor in determining whether such reactors could be deployed

quickly enough and at a large enough scale to make a significant contribution to reduction of carbon emissions by 2050, which is critical to mitigate the worst effects of climate change. The 2018 special report of the Intergovernmental Panel on Climate Change (IPCC) evaluated 85 primary energy supply scenarios to 2050 that would limit global mean temperature rise to 1.5°C (Rogelj et al. 2018). In these scenarios, the amount of nuclear power production in 2050 ranges from as little as two-thirds smaller than the 2017 level to nearly 13 times greater, with a median nuclear generation over all scenarios of about 2.5 times the 2017 level (Rogelj et al. 2018). This range reflects “both uncertainties in technological development and strategic mitigation portfolio choices” (Rogelj et al. 2018).

To achieve the IPCC report’s median projected increase of 150 percent in nuclear energy generation by 2050, nearly 600 GWe, or more than 500 large reactors, would have to be built worldwide—plus several hundred more that would be needed to replace reactors that will have reached the end of their operating lives. This would require that an average of at least two dozen reactors come on line each year between now and 2050. To put this in perspective, as mentioned above, the current rate of new reactor construction projects is below five new reactors per year—or only 20 percent of the rate corresponding to the IPCC’s median projection.

The IPCC report’s median increase in nuclear power deployment would thus be very challenging to achieve even with currently available LWRs. If the world must wait several decades for less mature NLWRs to become commercially available, it is hard to see how such reactors could be deployed quickly enough to play a significant role in limiting the worst impacts of climate change—even if they eventually turned out to be faster to build.

Is Development of NLWRs Essential for Nuclear Power’s Future?

As discussed above, while some observers argue that the future of nuclear power depends on development of NLWR designs, they have not made the case that the LWR has no future. Nearly all of the technological advances attributed to NLWR designs by the DOE and others (Petti et al. 2017) could also be realized in LWRs, including passive safety features, the potential for modular construction, the use of advanced fuels, greater plant autonomy to minimize labor costs, and underground siting. Indeed, some of these features have been incorporated into new LWR designs, such as the AP1000 and NuScale small modular reactor, although for economic reasons those reactors have other characteristics that may render them less safe than current-generation LWRs. What

is needed is a focused effort to develop LWRs that are genuinely safer *and* more economical at the same time.

There is also some potential for new types of LWRs to achieve one or both of the DOE's two strategic performance goals discussed above. The supercritical LWR, which is a Generation IV LWR with a coolant temperature of 500°C, could provide high-temperature process heat. And there are approaches for modestly increasing the sustainability of LWRs. Although significant research and development would be needed to achieve these goals safely and economically, commercializing NLWRs would introduce a no less difficult set of challenges. And thousands of reactor-years of operating experience gives the LWR an inherent advantage over even the more mature NLWR reactor concepts.

This report compares several classes of NLWRs to LWRs with regard to safety and security, the risks of nuclear proliferation and nuclear terrorism, and sustainability. Overall, the report finds little evidence that any of the NLWR designs currently under consideration, with the possible exception of once-through breed-and-burn reactors, would offer improvements over LWRs great enough to justify the expense, time, and risk necessary to commercialize and deploy them. Hence, one of this report's main conclusions is the bulk of nuclear energy-related research and development funding, both public and private, should be focused on improving the overall safety, security, efficiency, and cost-effectiveness of LWRs and the once-through fuel cycle.

Nuclear Power Basics

Many of the issues discussed in this report are technical and assume a familiarity with basic nuclear power concepts and terminology. This chapter provides background information that may be useful for understanding the technical analysis that follows.

The main objective of a nuclear fission power plant is to convert the energy released by the fission (the splitting) of atomic nuclei into electricity in a stable, reliable, and safe manner. (This report does not address nuclear fusion, another type of nuclear reaction that in principle could be used to generate electricity.) Nuclear fission reactions in a power reactor core generate heat, which is transferred to a system that converts it into electricity. Typically, a coolant, such as water, is circulated through the hot reactor core, and the heated coolant is then circulated through a power conversion system. For instance, heated coolant can be used to produce steam, which is then used to drive a turbine to produce electricity.

Given the laws of thermodynamics, not all of the heat generated by the reactor can be converted into electricity. A light-water reactor (LWR) that produces 3,300 MW of thermal energy (MW_{th}) would generate about 1000 MW of electricity (MW_e). The remaining heat energy is discharged to the environment as waste heat. Thus, less than one-third of an LWR's heat energy can be utilized to produce electricity. This fraction, known as the thermal efficiency, generally increases as the coolant temperature increases.

Nuclear Chain Reactions in the Reactor Core

When the nuclei of certain elements in the reactor fuel (called fissionable materials) are struck by neutrons, there is a chance

that they can undergo fission, releasing energy as well as additional neutrons. (Fissionable nuclei can undergo fission spontaneously; the likelihood of this occurring, which depends on the nature of the isotope, is always much lower than the likelihood of fission when struck by a neutron.) These neutrons can then strike other nuclei and potentially cause them to fission as well. A chain reaction can begin when at least one new neutron produced from fission is able to cause a second fission.

In any real-world reactor, some neutrons will be absorbed by fuel or other reactor materials, or even escape from the core, before they can induce a fission. When the average number of neutrons produced by fission is just enough to allow a self-sustaining chain reaction to occur, taking into account neutron losses from the system, then the reactor is said to be "critical." The power output of the reactor depends on the rate at which fission occurs, which in turn is related to the net number of neutrons in the core.

The likelihood that a neutron will interact with a nucleus generally increases as the neutron's speed decreases. Some nuclear reactors include materials called moderators that slow down fission neutrons, which can make it easier for fission to occur.

A fundamental aspect of any nuclear reactor is the arrangement of nuclear fuel, coolant, control rods, and (if needed) moderator materials in the core of the nuclear reactor so that the fuel can achieve a self-sustaining and stable neutron chain reaction. The coolant serves another critical function in addition to transferring the heat generated by the fissioning atoms, namely, ensuring that the temperature of the nuclear fuel remains at a safe level. If cooling is insufficient, the fuel can overheat, become damaged, and eventually melt, releasing highly radioactive materials into the environment.

The measure of how far the system is at any given moment from being exactly critical is its reactivity. A critical reactor has a reactivity of zero. If the reactor produces fewer neutrons than are lost, the reactor is subcritical, the reactivity is negative, and the power output will decrease. Conversely, if a reactor produces more neutrons than it loses, the reactor is supercritical, the reactivity is positive, and the power output will increase. If a reactor becomes supercritical, operators can insert control rods into the reactor to absorb neutrons and return it to a critical state. Conversely, if the reactor is subcritical, operators can withdraw the control rods. Controlling the chain reaction is essential to the safe operation of the reactor (see Box 2, p. 26).

The Components of Nuclear Fuels

To sustain a chain reaction and generate power, a nuclear fission reactor must be loaded with nuclear fuel appropriate for that reactor type. Nuclear fuels are composed of radioactive isotopes and other elements needed to make them chemically and mechanically stable under the harsh conditions of a reactor core.

FISSIONABLE, FISSILE, AND FERTILE ISOTOPES

A given element will always have the same number of protons, which is the atomic number of the element. For example, uranium (U) has an atomic number of 92. However, the number of neutrons can vary, and variants of an element with different numbers of neutrons are called isotopes. The different isotopes of an element are identified by their total number of protons and neutrons. Some of the important isotopes of uranium are U-235 and U-238.

Fissionable isotopes are those capable of being split when struck by a neutron. Some fissionable isotopes are also fissile: they can be fissioned by neutrons of any energy, including low-energy (“thermal”) neutrons. Other fissionable isotopes can only undergo fission if struck by a neutron with an energy above some minimum value; however, when struck by low-energy neutrons, they can be transmuted into fissile isotopes such as plutonium-239. These are called fertile isotopes. There is also a chance that a fissile nucleus will not fission when struck by a neutron, but instead will capture it and transmute into a heavier isotope. The relative likelihood that a neutron will fission a nucleus versus be absorbed by it is the fission-to-capture ratio.

The term “fissile material” is commonly used to denote nuclear materials that can be used to make nuclear weapons. Confusingly, this has a different meaning than “fissile isotope.” In fact, some fissile materials are mixtures of fissile and fertile isotopes. “Nuclear explosive material” or simply “weapon-

usable material” are preferable, more precise terms for nuclear materials that can be used directly to make nuclear weapons.

FISSION PRODUCTS

When a nucleus undergoes fission, in addition to releasing neutrons and energy, it splits into other nuclei called “fission products,” some of which are intensely radioactive. An operating reactor core will contain hundreds of different fission product isotopes with a wide range of different half-lives—the period of time after which half of a given quantity of radioactive material will have changed into other isotopes, known as “decay products.” Two key fission products important to nuclear safety include iodine-131, with a half-life of eight days, and cesium-137, with a half-life of 30 years. The presence of cesium-137 is one of the main reasons why spent reactor fuel emits a very hazardous radiation field and must be handled remotely for many decades after discharge.

THE ELEMENTS IN NUCLEAR FUEL

Depending on the reactor design, various combinations of isotopes of uranium, plutonium, thorium, or other elements (neptunium, americium, curium) can be used in nuclear fuel. These elements are known as actinides.

URANIUM (U)

Uranium (U) is the element most commonly used for nuclear reactor fuel; it is categorized by the relative amounts of the isotopes U-235 and U-238 it contains.

Natural uranium. Natural uranium ore is primarily composed of two isotopes of uranium: approximately 99.3 percent U-238 (fertile) and 0.7 percent U-235 (fissile). Some types of reactors can use natural uranium as fuel, but they require moderators other than ordinary water, such as graphite.

Enriched uranium. Most nuclear power reactors operating today must use enriched uranium, a fuel with a higher concentration of U-235 than natural uranium. Uranium enrichment is a complex and expensive process. (Although civil enrichment facilities are configured for optimal production of low-enriched uranium for reactor fuel, they pose nuclear proliferation risks because they can be readily modified to enrich uranium to the higher levels needed for use in nuclear weapons.)

The various grades of enriched uranium include the following:

Low-enriched uranium. Low-enriched uranium (LEU) is enriched to a U-235 concentration greater than 0.7 percent (the concentration in natural uranium) and below 20 percent. LWRs, which use ordinary water as a coolant and moderator,

continued on p. 27

Reactor Stability: Controlling the Chain Reaction

The stability of a reactor's power output is a critical aspect of nuclear safety. When a reactor operates, small changes in its state occur constantly, which can affect the rate of nuclear fission and hence the amount of power produced. A well-designed reactor will respond to those changes in a slow and predictable manner, providing ample time for operators to take corrective actions.

The reactivity feedbacks of a reactor are measures of how it will respond to a change in operating conditions—ranging from small fluctuations in temperature to major events such as a loss of cooling due to a pipe break. Will the disturbance cause the fission rate and power level to decrease or increase—and how quickly?

COEFFICIENTS OF REACTIVITY

The overall reactivity of a reactor—and its response to changing conditions and consequent stability—will depend on numerous factors, including its physical size, the temperature of the coolant, the fuel, the moderator (if there is one), and reactor structural elements. The effects of these various factors are described by coefficients of reactivity. A positive coefficient indicates that an increase in a parameter (such as temperature) will increase the reactivity—thus creating a positive feedback loop where the reactor power will increase, further increasing the temperature—a potentially unstable condition.

The inherent stability—and therefore safety—of a reactor depends on how the reactivity of the system will respond to changes without intervention by the operator or the activation of automatic control systems that will not always work.

Some of the important coefficients are:

- The **moderator temperature coefficient of reactivity**, which indicates how the reactor will respond to a change in temperature of the moderator, if one is present.
- The **coolant temperature coefficient of reactivity**, which indicates how the reactor will respond to a change in temperature of the coolant. In an LWR, the coolant—light water—is also the moderator.
- The **fuel temperature coefficient of reactivity**, which indicates how the reactor will respond to a change in the temperature of the fuel. This coefficient is referred to as “prompt” because the fuel temperature responds almost immediately to an increase in power, whereas the coolant

and moderator temperatures take a few seconds to adjust. Thus, a negative fuel temperature coefficient is a critical component of the reactor's overall inherent reactivity feedback.

An important phenomenon that affects the fuel temperature coefficient is known as the Doppler effect. As the fuel temperature increases, non-fissile nuclei such as U-238 in the fuel absorb more neutrons but do not fission, thus slowing down the fission process. This feedback is nearly instantaneous, because it is a response to the heating of the fuel itself. The magnitude of the effect depends on the properties of the fuel and the neutron speeds within the reactor.

- The **void coefficient of reactivity**, which indicates how the reactor will respond to changes in the number and size of bubbles, or voids, that appear or expand in the coolant as it heats up. These voids are regions where the density of the coolant—and therefore its neutron-moderating and -absorbing effects—is greatly reduced.

A reactor with all negative coefficients will experience only a minimal rise in power production and temperature if a change in operating conditions causes an unplanned increase in reactivity and therefore the rate of fission. This behavior enhances the reactor's stability. For reactors with a mix of positive and negative coefficients, the situation is more complicated, and the overall reactivity of the reactor is calculated using information from computer modeling and experiments. These calculations sometimes have large uncertainties, making it difficult to accurately assess the reactor's stability.

REACTOR STABILITY THROUGH DELAYED NEUTRONS

Another important factor in determining reactor stability is the presence of “delayed” neutrons. Most of the neutrons in a reactor are prompt, or generated immediately after fission, but a small fraction are delayed—emitted by certain fission products up to nearly a minute after fission occurs. These delayed neutrons increase the time scale over which reactivity changes in response to perturbations of the system. Reactors with a larger fraction of delayed neutrons are more stable, responding more slowly to reactivity perturbations and providing more time for control of the chain reaction.

continued from p. 25

use LEU fuel, typically with enrichments of 3 to 5 percent U-235. Because it is impractical or even impossible to use LEU directly in a nuclear weapon, depending on the enrichment, it poses far lower nuclear proliferation and nuclear terrorism risks than highly enriched uranium.

Highly enriched uranium. Highly enriched uranium (HEU) is enriched to a U-235 concentration of 20 percent or more. While some power reactors historically have used HEU fuel, such use has been discouraged in recent decades because of its proliferation potential. HEU at any enrichment can be used to make nuclear weapons, with more material required for lower enrichment levels. However, HEU used in nuclear weapons is typically enriched to 90 percent U-235 or greater.

High-assay low-enriched uranium. High-assay LEU (HALEU) is a sub-category of LEU with a U-235 enrichment at or above 10 percent and below 20 percent.⁸ HALEU of various enrichments would be needed for some NLWR reactor designs. Although HALEU is a type of LEU, it poses greater proliferation and security risks than the lower-assay LEU used in LWRs. Some experts have said that HALEU can be used to make nuclear weapons, although with much greater difficulty than with HEU (Mark 1984). In accordance with this higher risk, HALEU requires greater security than LEU with enrichments below 10 percent.

LEU+. LEU+ is LEU with an enrichment greater than 5 percent but less than 10 percent. This enrichment range is being considered for use in certain types of new fuels for LWRs. It is classified as having the same security risk as LEU with enrichments of 5 percent or below.

Depleted uranium. Depleted uranium (DU) has a U-235 content of 0.3 percent or below. It is a byproduct of the uranium enrichment process. It cannot be used by itself as nuclear reactor fuel and is generally considered a waste product.

The amount of effort needed to enrich a given quantity of uranium to a specified U-235 concentration is called separative work, measured in separative work units (SWU). For example, starting with natural uranium, it takes roughly 30 times as much SWU to produce 1 kilogram of 90 percent-enriched HEU than the same quantity of 4.5 percent-enriched LEU. However, if one starts with 4.5 percent LEU, it would only take about one-third as much SWU to produce 1 kg of 90 percent HEU than if one started with natural uranium.

PLUTONIUM (PU)

Plutonium does not exist naturally but is produced in nuclear reactors when uranium fuel is irradiated. When the fertile isotope U-238 captures a neutron, it undergoes two radioactive

decays and is transmuted to Pu-239. (Further neutron capture will produce higher isotopes of plutonium.) Plutonium, like HEU, is a nuclear explosive material. In contrast to uranium, all isotopic combinations of plutonium can be used to build nuclear weapons (except for pure Pu-238, which generates decay heat at a high rate, making it impractical in a weapon). Plutonium can also be used as fresh fuel for reactors, but such fuel poses greater proliferation and terrorism risks than LEU fuel.

LWR spent fuel contains about 1 percent Pu by weight. To extract and concentrate plutonium for reactor fuel or for weapons, spent fuel must be reprocessed. The plutonium in LWR spent fuel is diluted and embedded in large, heavy, and highly radioactive spent fuel assemblies, making recovery difficult. Therefore, reprocessing is a complex and challenging process. From a chemistry perspective, reprocessing is somewhat easier than uranium enrichment because it involves separating different elements rather than different isotopes of the same element. However, since spent fuel is highly radioactive, it can only be reprocessed in heavily shielded facilities utilizing remote-handling equipment. But if the end product—plutonium—is successfully separated, it is not highly radioactive, and a weapon's worth of material—less than 10 kilograms—can be easily carried by a single person.

For this reason, nuclear fuel cycles that separate plutonium for reuse as reactor fuel have inherent security and proliferation risks because they greatly increase the vulnerability of plutonium to theft or diversion. (Such fuel cycles are referred to as closed, whereas those that dispose of the spent fuel directly are considered open.)

OTHER TRANSURANIC ELEMENTS (TRU)

In the same way that plutonium-239 is produced when U-238 absorbs a neutron, successive neutron capture will produce elements with higher atomic numbers than that of uranium, which is 92. Such elements are referred to as transuranic elements (TRU). Plutonium, with an atomic number of 94, is a transuranic element. Other transuranic elements, also referred to as minor actinides, are neptunium (Np), americium (Am), and curium (Cm), with atomic numbers of 93, 95, and 96, respectively.

Several Np, Am, and Cm isotopes can be used as fuel for nuclear weapons, although doing so is generally more technically difficult than using plutonium.

Transuranic elements other than plutonium are not useful as fuel in LWRs because they have a greater tendency to absorb thermal neutrons and transmute into heavier isotopes than undergo fission and release energy. However, they can more effectively be fissioned by fast neutrons and can be used as fuel for fast reactors.

As with plutonium, in order to use Np, Am, or Cm for reactor fuel or for nuclear weapons, they first must be separated from spent fuel by reprocessing. However, since they are present in LWR spent fuel at lower concentrations than that of plutonium, they are more difficult to recover. More spent fuel would have to be reprocessed to obtain quantities useful for either of those purposes.

THORIUM

Thorium (Th), with an atomic number of 90, is an actinide element like uranium. It is believed to be at least three times more abundant than uranium in the Earth's crust, although that estimate has recently been questioned (NEA 2015). Natural thorium is almost entirely composed of the isotope Th-232, which is not fissile but is fertile. Natural thorium cannot be enriched like natural uranium; therefore, in order to use thorium as reactor fuel, it must be mixed with fissile isotopes such as U-235 or U-233.

When fertile Th-232 is irradiated, it can capture a neutron and transmute into the fissile isotope U-233 through the intermediate product protactinium-233 (Pa-233), which has a half-life of 27 days. Thus Th-232 can be converted to a usable fuel isotope in a similar manner to the transmutation of U-238 to Pu-239, although the half-life of the intermediate decay product poses an additional complication, as discussed in chapter 7.

REACTOR FUEL MATERIALS

The fuel into which fissionable materials are incorporated must be physically and chemically suitable for the harsh conditions of reactor operation. The conventional fuel for LWRs is a uranium oxide ceramic, formed into pellets and stacked in long, thin metal tubes known as cladding. These fuel rods are bundled into assemblies. The cladding material is typically an alloy containing the metal zirconium known as Zircaloy.

To use plutonium or other TRU in reactor fuel, they must be blended with uranium prior to being formed into fuel pellets (or whatever form the fuel takes). One example is mixed-oxide fuel for LWRs, a blend of plutonium and uranium oxides. Mixed-oxide fuel is a less attractive fuel for LWRs than LEU fuel because it is more expensive and requires more stringent security measures.

Besides oxides, various types of reactors can use metal, carbide, or nitride fuels. Some types of reactors can even use liquid fuels, such as molten salts. The design of such reactors is quite different from those using solid fuels, as discussed below.

Thermal and Fast Reactors

Nuclear reactors have two main variants: thermal reactors and fast reactors. These terms refer to the average speed of the neutrons in the reactor. The major difference is that thermal reactors have moderator materials that significantly slow down the neutrons, whereas fast reactors do not. Reactors that use coolants other than water can be either fast or thermal reactors, depending on the properties of the coolant and other design features.

Thermal reactors use a moderator such as “light” (or ordinary) water (which also serves as the coolant) because fuel nuclei have a much higher chance of interacting with slower neutrons and undergoing fission than with faster ones. Because the probability that these isotopes will fission is greater in thermal reactors, the fuel can have a relatively low concentration of fissile material, such as the LEU fuel enriched to 3 to 5 percent that is used in LWRs.

To compensate for a lower probability of fission, fast reactors must use fuel with a higher concentration of fissile material—historically either HEU or a mixture of uranium and at least 12 to 15 percent plutonium. Such reactors pose security concerns because HEU and plutonium can be used directly to make nuclear weapons. As discussed in chapter 4, the fresh fuel must therefore be stringently secured. Because of the security risks of plutonium and HEU, some proposed fast reactors are being designed to use HALEU, although, as discussed above, HALEU fuel also requires greater security than the LEU fuel that LWRs use.

BREEDING

Given fast reactors' need for fuels that require greater—and more costly—security than LWR fuel, why would anyone build one? The historical motivation is that a fast reactor, in theory, can “breed”—that is, it can generate as much or even more fissile fuel than it consumes.

As scientists learned early in the development of nuclear power, when isotopes are fissioned by fast neutrons, they release greater numbers of neutrons on average than when fissioned by thermal neutrons. These extra neutrons are not needed to maintain the nuclear chain reaction that keeps the reactor operating. Instead, they can be used to convert fertile U-238 into fissile Pu-239.

In most thermal reactors, not enough extra neutrons are generated to breed new fuel. (One exception is the thorium-fueled MSR, discussed in chapter 7). In contrast, in a fast reactor there are enough excess neutrons to breed plutonium. The extra plutonium can then be used to refuel the reactor

and even provide fuel for a new reactor, if enough is bred. Typically, the reactor core includes both “driver” fuel containing plutonium and “blanket” fuel containing U-238 (in the form of natural or depleted uranium). The driver fuel is the main source of heat that can be used to generate electricity, while most of the breeding occurs in the blanket fuel.

In conventional fast reactors, for the plutonium in the blankets to be used as new fuel, the blankets must be reprocessed to extract the plutonium from the remaining U-238. The residual plutonium in the spent driver fuel—which is significant—is also recovered by reprocessing. The recovered plutonium and U-238 can then be used in fresh driver and blanket fuel.

The potential for nuclear reactors to generate their own fuel was initially seen as an essential feature in the early days after the Manhattan Project, when uranium was thought to be scarce—and what was available was reserved for the nuclear weapons program. However, that rationale is much less compelling today now that uranium has proven to be an abundant natural resource.

BURNING

Another way that fast and thermal neutrons differ is their propensity to induce fission when striking certain isotopes, rather than to simply be absorbed. As discussed above, successive neutron capture will produce TRU isotopes in a reactor core. If TRU isotopes such as Pu-240 or americium-241 are struck by a thermal neutron, they have a high probability of absorbing the neutron and transmuting into a heavier isotope. However, if they are struck by fast neutrons, they are more likely to fission. Fast reactors can therefore use TRU isotopes as fuel far more effectively than thermal reactors. Since many of these TRU isotopes are long-lived and generate significant decay heat, they could potentially cause problems for nuclear waste disposal. Thus the ability of fast reactors to more efficiently fission TRU isotopes is often cited as an advantage over thermal reactors. Some observers refer to this as nuclear waste “burning,” even though the TRU elements are only a small component of the total mass of nuclear waste. This is discussed in more detail in chapter 3.

Fuel Burnup and Refueling

As noted above, fresh LWR fuel is typically composed of low-enriched uranium oxide, containing 3 to 5 percent U-235, with the balance being primarily the isotope U-238. As the fuel is irradiated in the reactor, it undergoes both changes in radionuclide composition and changes in its physical and chemical form.

In LWRs, which have a thermal neutron spectrum, fission is much more likely to occur in U-235 than in U-238, and most of the energy production is due to fission of U-235. While U-238 is much more likely to absorb a neutron than fission in a thermal reactor, when it does so it can be transmuted to Pu-239, which is more likely to fission than absorb a neutron.

Thus, as the fuel is irradiated, energy is released by fission of both U-235 and Pu-239. The initial amount of U-235 is depleted as it undergoes fission. This is compensated for to some extent by conversion of U-238 to Pu-239. Eventually, however, the amount of U-235 plus Pu-239 becomes too low to sustain the nuclear chain reaction and the fuel becomes spent (no longer usable). For this reason, nuclear fuel can only be used for a limited time before it must be discharged from a reactor core and replaced with fresh fuel.

A second limiting factor for how long nuclear fuel can be used is the degradation of fuel matrix and cladding materials as they are subject to high heat, chemical interactions, exposure to radiation, and pressure from fission product gases. Eventually, the fuel becomes so degraded that it cannot safely remain in the reactor without risk of rupture.

The “fuel burnup” is a measure of the amount of heat (usually expressed in terms of megawatt-days of thermal energy, or MWD—the “thermal” is implied) generated by the irradiation of one metric ton of heavy metal (MTHM) in the fresh fuel—that is, the initial quantity of uranium and other “heavy” fissionable materials such as plutonium. The burnup depends on both the power density—how much of the reactor’s power is generated by a given quantity of fuel—and the length of time that the fuel remains in the reactor core.

If a metric ton of uranium could be fissioned completely—that is, 100 percent burnup—it would release about 970,000 MWD of thermal energy. However, it is not possible to achieve such a high burnup in a realistic reactor system. A typical average discharge burnup for LWR fuel is 50,000 MWD/MTHM, which corresponds to fission of around 5 percent (50,000/970,000) of the initial uranium content. Spent fuel discharged from an LWR at this burnup contains less than one weight percent U-235 (compared to 4 to 5 percent in the initial fuel), and just over one weight percent total plutonium. The spent fuel also contains about 0.1 weight percent of other TRU, such as americium-241. Fission products make up about 5 percent by weight. The balance, around 93 weight percent, is almost entirely U-238.

The length of time between reactor refueling outages is related to the peak allowable burnup of the fuel. Typically, a reactor core will have several batches of fuel that were loaded in the reactor at different times. During refueling, the oldest fuel is removed, and the remaining fuel is shuffled to ensure

that the core generates power evenly and that safety margins are maintained. Generally, irradiating fuel to higher burnup enables the fuel to remain in the reactor longer and allows for less frequent refueling.

Reactor Safety Considerations

Nuclear power plants generate high rates of heat and produce large quantities of highly radioactive materials—a potentially dangerous combination. If a reactor core generates heat at a higher rate than the coolant system is able to remove it, the fuel and reactor structures can be damaged and a catastrophic release of radioactivity can occur.

TYPES OF ACCIDENTS

The terms “design-basis” accidents and “beyond-design-basis” accidents are commonly used. Design-basis accidents are those that are taken into account in the design of the reactor. Safety systems are provided to protect against design-basis accidents and prevent them from causing large releases of radioactivity. Historically, “beyond-design-basis” (also called “severe”) accidents have been considered to be less probable than design-basis accidents, although they can and have occurred. Beyond-design-basis accidents can overwhelm safety systems, leading to a core melt and large radioactivity release.

Most initiating events that can trigger beyond-design-basis accidents and core meltdowns at nuclear reactors can be classified in three types: (1) a rapid increase in the rate of nuclear fission (that is, an increase in reactivity) and an uncontrollable increase in power; (2) a loss of coolant due to leakage or inadequate coolant flow, causing the reactor fuel to overheat; and (3) a loss of the ability to remove heat from the reactor system (such as the total loss of electric power—i.e., a station blackout), which could also lead to core melt.

The three most serious nuclear power plant accidents—all of which could be considered “beyond-design-basis”—illustrate these three categories. The Chernobyl Unit 4 explosion in 1986 in the former Soviet Union was initiated by a rapid increase in reactivity. The Three Mile Island Unit 1 meltdown in 1979 in Pennsylvania was a loss-of-coolant accident caused by a stuck-open valve. And the Fukushima Daiichi accident in 2011 in Japan, which caused three reactors to melt down and release radioactivity, resulted from a loss of heat removal triggered by a loss of the electrical power needed to operate coolant pumps and other safety systems.

Accident initiators can be further classified into two types: internal events that stem from problems occurring within the nuclear plant, and external events that are

triggered by natural disasters and other types of incidents originating outside the plant. The Chernobyl and Three Mile Island accidents were caused by internal events (including operator errors), whereas the total loss of electrical power at Fukushima had an external cause—a severe earthquake that took down power lines and site flooding from subsequent tsunamis that damaged electrical generating and distribution equipment. Intentional acts—known as acts of “radiological sabotage”—can also be accident triggers, by initiating conditions similar to internal and/or external events. Indeed, knowledgeable saboteurs could quickly induce conditions resulting in core damage and radiological releases that would be highly unlikely to occur solely by chance.

SAFETY SYSTEMS

Current-generation reactors typically have multiple backup safety systems to protect the reactor in the event of an accident. They also have several physical barriers to prevent the escape of radioactivity into the environment in the event that the fuel is damaged, including a metal vessel surrounding the fuel and a leak-tight containment structure made of steel and concrete. Another layer of safety consists of pre-planned actions to protect the public, such as evacuation or sheltering, within an emergency planning zone around the reactor. In addition, the United States and some other countries require that nuclear plants have armed security personnel to protect against radiological sabotage. These diverse and redundant safety and security measures are referred to as “defense-in-depth.”

However, accidents or acts of sabotage can be severe enough to disable multiple safety systems, making core melt inevitable. When that occurs, the nuclear fuel heats up to a temperature at which it begins to degrade and eventually melt, releasing radioactive fission products into the coolant system. The excess heat also increases the temperature and pressure within the reactor and containment structure. Eventually, the hot molten core will slump to the floor of the reactor vessel and melt its way through into the containment structure. The increases in temperature and pressure, as well as explosions of combustible gases such as hydrogen, can cause the containment to fail, releasing radioactivity into the environment.

How severe such a release could be for public health and the environment is largely determined by the “source term”—the types of isotopes that are released, their quantities, their chemical forms, and other factors relevant to how the materials are released and dispersed. In addition, prevailing weather conditions and the population distribution in the vicinity of the reactor are important factors.

With the wide range of variables involved, it is difficult to develop a simple way to compare the overall safety of different reactor types. To rigorously determine whether any advanced reactor would be safer overall than current-generation LWRs, one would need to sum up the risk of a large radiological release over all potential severe accident sequences, including waste storage accidents, and compare it to the risk associated with a current-generation LWR. This would require a comprehensive probabilistic risk assessment, validated with data from operating experience.

While probabilistic risk assessments for LWRs have operating experience to draw upon for validation, achieving the same level of validation remains far in the future for any NLWR design. And even the best risk assessments have large uncertainties associated with unknowns such as the risks of catastrophic external events, human errors, and sabotage. Thus, qualitative safety measures such as defense-in-depth, which are needed to compensate for such uncertainties, need to be given great weight in comparative assessments.

Nuclear Power Sustainability

The operating light-water reactor (LWR) fleet uses mined uranium for fuel and generates highly radioactive nuclear wastes. Both the front and the back ends of this fuel cycle have the potential for significant health and environmental impacts if not rigorously managed.

Two Primary Goals for Increasing Sustainability

One of the primary goals cited by NLWR developers is to reduce these impacts by increasing the “sustainability” of nuclear power—or, as the Department of Energy (DOE) puts it, to “extend natural resource utilization” and “reduce the burden of nuclear waste for future generations” (Petti et al. 2017). In other words, for a nuclear reactor to be more sustainable than an LWR, it should (1) use natural uranium more efficiently than LWRs, and (2) generate less nuclear waste requiring long-term disposal—or even use fuel obtained by reprocessing and “recycling” the mountain of highly radioactive nuclear waste that LWRs have already produced, more than 80,000 metric tons and counting in the United States today.

However, although these goals certainly sound worthwhile, it is not clear whether achieving them is practical or even necessary for the future of nuclear power. Two fundamental questions need to be addressed. First, to what extent would any NLWR and its associated fuel cycle be significantly more sustainable in practice than the LWR once-through cycle? And second, would those benefits be significant enough to justify the substantial investment required to develop and deploy such a reactor at a large scale? These highly complex

questions depend on many variables and are very sensitive to model assumptions. While it is beyond the scope of this report to fully answer these questions, this chapter discusses key issues that must be considered.

REDUCED LEVELS OF LONG-LIVED RADIOACTIVE WASTE

Spent fuel from LWRs contains highly radioactive, long-lived isotopes that must be isolated from the environment for hundreds of thousands of years to protect public health and the environment. The only way this can plausibly be achieved is to dispose of the waste in a robust underground facility known as a geologic repository. However, most countries with nuclear plants, including the United States, have failed to open geologic repositories for spent nuclear fuel. Only Finland, with a much smaller amount of nuclear waste than the United States, is making steady progress.

Highly radioactive wastes requiring disposal in a deep geologic repository are generated by all reactors and fuel cycles. However, some advocates of reprocessing argue that, given the political difficulties and technical challenges of establishing repositories, geologic disposal space will be scarce and valuable in the future and must be conserved by reducing nuclear waste volume (Bailey 2014). A new reactor design could reduce the future waste burden if it produced less long-lived waste than an LWR while generating the same amount of electricity. Furthermore, if the reactor could efficiently use actinides extracted from existing LWR waste as new fuel—often misleadingly referred to as “burning” nuclear waste—this approach could reduce the repository space needed for the current waste stockpile.

EFFICIENT USE OF URANIUM

While reducing the amount of uranium used by nuclear reactors to generate a given amount of energy could conserve uranium resources, uranium is currently not in short supply; therefore, there is no economic driver at present for such a change. Early in the nuclear era, estimates of worldwide uranium ore were low, and the nuclear power community feared that there would not be enough uranium to fuel reactors in the future. But these estimates have risen over time, and there is little risk that the world will run out of uranium for the foreseeable future.

The latest assessment of resources by the Nuclear Energy Agency and the International Atomic Energy Agency in 2020 found that identified recoverable uranium resources would be sufficient to fuel the global nuclear reactor fleet for more than 135 years at the 2019 rate of consumption (just under 400 gigawatts of electricity) (NEA 2020). Better recovery methods could make available up an additional 40 years' worth of consumption. Thus, even if nuclear energy generation worldwide were to double over the next few decades—more than the projected 80 increase in the International Atomic Energy Agency's current “high case” scenario for growth by 2050—identified resources would likely be adequate until the end of the century. In addition, sources of uranium that are believed to exist but remain undiscovered are estimated to be nearly as great the currently identified resources (NEA 2020). And ultimately, the world's oceans, which contain a vast quantity of uranium at a low concentration, serve as a backstop to any supply shortage. Although the cost of uranium will increase as more readily exploitable resources are depleted, that should be compared to the additional costs and risks associated with developing and operating new reactor types that are more uranium-efficient.

However, resource depletion is not the only concern associated with uranium consumption. Uranium mining is dangerous for workers and pollutes soil, air, and groundwater. Uranium mining is less widespread in the United States today than in the past, but over time it has left thousands of abandoned mines and dozens of uranium processing sites that require cleanup, many of which are located within the Navajo Nation and continue to have a disproportionate impact on the Navajo people. Moreover, uranium waste dumps and mines emit carcinogenic radon gas decay products that pose health risks to both miners and individuals living downwind. More modern mining and processing methods, although less damaging than historical practices, can also harm public health and the environment if not implemented with the most rigorous standards and oversight.

Reactors that use uranium more efficiently could have health and environmental benefits by reducing the need for mining. However, the benefits from reducing uranium mining activities would have to be balanced against the increased environmental risks of more uranium-efficient reactors and their fuel cycles. Increasing uranium efficiency usually entails reprocessing spent fuel, which generates a number of different radioactive waste streams and emits radioactive gases into the atmosphere—many with wide-reaching health and environmental impacts themselves.

To maximize the utilization of natural uranium, NLWRs would have to be capable of effectively using depleted uranium—the leftover material produced during enrichment—as fuel. Depleted uranium has a U-235 content of 0.3 percent or below. Only a small fraction of mined natural uranium ends up in the enriched uranium fuel used in LWRs; the depleted uranium “tails” of the process are stored as waste requiring disposal. The production of one year's supply of enriched uranium for a typical LWR—20 metric tons—generates about 180 metric tons of depleted uranium. This material has accumulated as waste in the United States and most other countries because it is not economical today to re-enrich it for use as LWR fuel. The DOE now holds more than 500,000 metric tons of uranium tails in the form of uranium hexafluoride gas, requiring hundreds of football fields' worth of storage space. Although this material poses a relatively low radiological hazard in storage, it will likely require disposal in a deep geologic repository in the long term, but there is no clear disposition path at present.

The Challenging and Conflicting Goals of Sustainability

Many NLWR developers argue that their systems will achieve breakthroughs in improving nuclear power sustainability. A good example is the Argonne National Laboratory, a DOE facility, which has been developing sodium-cooled fast reactor technology (see chapter 1) and an associated fuel reprocessing system (known as pyroprocessing) for decades. In a 2012 brochure, ANL claimed that its pyroprocessing technology, used in conjunction with fast reactors, would turn nuclear waste into a “wonderfuel” (ANL 2018).

Specifically, Argonne National Laboratory asserts that its fast reactor and pyroprocessing system would:

- “allow 100 times more of the energy in uranium ore to be used to produce electricity compared to current commercial reactors”
- “ensure almost inexhaustible supplies of low-cost uranium resources”

- “markedly reduce the amount of waste and the time it must be isolated—from approximately 300,000 to approximately 300 years—by recycling all actinides” (ANL 2018)

The first two bullets refer to increasing uranium efficiency and the third to reducing the waste disposal burden.

While this reactor system certainly sounds promising, this study finds that these claims are highly misleading. First, it is important to note that these two aspects of sustainability—significantly reducing the quantity of TRU elements (primarily neptunium, plutonium, americium, and curium) contained in nuclear waste and significantly increasing uranium utilization efficiency—cannot be *simultaneously* achieved with the same reactor and fuel cycle system. The two goals are technically incompatible. This is because a nuclear reactor can only extract energy from a fixed amount of fissionable material per year, which depends on its power level. If the energy is produced by the fission of a TRU element that comes from nuclear waste, it cannot be produced by the fission of new fissionable materials generated from depleted uranium.

To significantly reduce the high-level waste disposal burden, the reactor system would be designed to prioritize the fission of long-lived TRU isotopes extracted from nuclear waste—thus, most of the energy it produces would result from TRU fission. That is, the TRU contained in the LWR spent fuel stockpile would be the primary makeup source of fissionable material for fresh fuel. On the other hand, to significantly increase the efficiency of uranium utilization, as discussed above, a reactor system must produce most of its energy by converting the U-238 in the depleted uranium to plutonium and then fissioning the plutonium. In this case, the depleted uranium stockpile would be the primary source of fresh fuel. But because the amount of energy produced per year in a reactor is constant, it cannot effectively use the existing stockpile of TRU in nuclear waste and the existing stockpile of depleted uranium at the same time.

Moreover, while attaining either sustainability goal individually may be achievable on paper, neither can be attained in practice over a reasonable time scale, as both would require a level of system performance far beyond what nuclear facilities are capable of today or are likely to achieve in the foreseeable future. In order to make good decisions regarding the development of reactors systems with greater sustainability, it is critical that expectations for their real-world performance be distinguished from their theoretical performance in an ideal world.

High-Level Waste Reduction

The United States has a nuclear waste problem—as do almost all other nations with nuclear power plants. Today, no country has a geologic repository ready to accept spent fuel or high-level waste, and only Finland is constructing one for a nuclear power sector much smaller than that of the United States or other larger countries. While the United States does operate an underground repository—the Waste Isolation Pilot Plant in New Mexico—the facility accepts only TRU-containing wastes from military activities. It is legally prohibited from accepting spent fuel or high-level radioactive waste.

Decades ago, the United States decided as a matter of policy to dispose of its spent fuel and high-level waste in a deep underground mined repository. However, for political, technical, and legal reasons, it has not yet been able to successfully build such a repository. It officially chose Yucca Mountain in Nevada as the repository site in 2002, and the DOE applied to the Nuclear Regulatory Commission for a construction license in 2008. Two years later, the DOE withdrew its application, stating that Yucca Mountain was not workable. Although early in former President Donald Trump’s tenure the DOE attempted to provide funds to restart Yucca Mountain project licensing at the Nuclear Regulatory Commission, the requests were rebuffed by Congress. In its budget request for fiscal year 2021, the DOE did not seek funding to move Yucca Mountain forward.

Nevertheless, the site remains the only one in the United States designated by law for geologic disposal of spent nuclear fuel. In order to prevent the disposal burden from being imposed on only one state, the law currently limits the capacity of Yucca Mountain to 70,000 metric tons heavy metal (MTHM) of waste. The US stockpile of waste has already exceeded this limit—as of mid-2017, US commercial reactor sites stored nearly 80,000 MTHM of spent fuel (GAO 2019). Subsequently, the reactor fleet has added about 2,000 MTHM of waste per year to this stockpile.

However, the physical capacity of Yucca Mountain is four to nine times greater than maximum amount of waste it is legally allowed to store (Maden 2009). If the statutory limit were relaxed or eliminated, the United States might not need a second repository for centuries. A bill passed by the House of Representatives in May 2018 would increase the capacity to 110,000 MTHM, and the bill was introduced in both houses of Congress in the 2019-2020 session, but no votes were taken by either house.

CONSUMING NUCLEAR WASTE: THE HOPES AND CLAIMS

Given the lack of progress on spent fuel disposal, the notion that “advanced” nuclear reactors could consume existing nuclear waste is a very compelling idea to the public and to many policymakers. For example, Senator Sheldon Whitehouse (D-RI), in a floor speech in February 2016, referred to “advanced reactors that could actually consume spent fuel from conventional reactors and help us draw down our nuclear waste stockpile” as “the Holy Grail” (Whitehouse 2016).

Senator Whitehouse should not be faulted for his enthusiasm—many credible nuclear experts make assertions that NLWRs could essentially eliminate nuclear waste. For instance, the American Nuclear Society’s statement “Fast Reactor Technology: A Path to Long-Term Energy Sustainability” states that for a fuel cycle with fast reactors and reprocessing, “virtually all long-lived heavy elements are eliminated during fast reactor operation, leaving a small amount of fission product waste that requires assured isolation from the environment for less than 500 years” (ANS 2005). This statement is echoed by a number of NLWR developers who say that their designs could “consume,” “burn,” or “recycle” spent fuel from LWRs. These include not only liquid metal-cooled fast reactors, but also gas-cooled reactors and molten salt reactors. One example—Argonne National Laboratory—has been cited above. Other examples follow.

OKLO, INC.’S FAST MICROREACTOR

Jacob DeWitte, co-founder of Oklo, Inc, testified before Congress that the company’s 4 megawatt-thermal fast micro-reactor, now called the Aurora and under licensing review by the NRC, “can consume the used fuel from today’s reactors” (DeWitte 2016).

GENERAL ATOMICS’ ENERGY MULTIPLIER MODULE

General Atomics has been developing a high-temperature gas-cooled fast reactor—the Energy Multiplier Module (EM²), with a power output of 265 megawatts-electric (MWe). According to General Atomics, “deployed in sufficient numbers, EM2 is capable of substantially reducing pressures for long-term storage and turning our waste stockpile into an important energy resource” (GA 2019).

SEABORG TECHNOLOGIES’ COMPACT MSR

Seaborg Technologies, a Denmark-based company is developing a thorium MSR that it has also referred to as a “waste-burner” (Seaborg Technologies 2015). The company says that “realizing the waste burning potential is part of Seaborg’s mission to make nuclear truly sustainable” (Seaborg Technologies n.d.).

TRANSATOMIC POWER’S WASTE-ANNIHILATING MSR

Another MSR startup, Transatomic Power had claimed that its reactor could consume nuclear waste as fuel. However, after errors were discovered in its analyses, it had to backtrack on the claim and lost credibility before shutting down in September 2018 (see chapter 7).

CONSUMING NUCLEAR WASTE: THE REALITY

The story of Transatomic Power is a cautionary tale for other NLWR developers who overstate the nuclear waste burning capabilities of their reactor systems. Unfortunately, as discussed below, it is virtually impossible to completely eliminate or even significantly reduce nuclear waste by using it as fuel in any real reactor system. Therefore, the United States will need a deep geologic repository for nuclear waste regardless of the types of reactors it uses in the future.

What most NLWR developers actually mean by “consuming” nuclear waste is using some of the components of spent fuel—namely, plutonium and other fissionable TRU isotopes—as fresh fuel in their reactors. These isotopes have half-lives from hundreds to millions of years. When they undergo fission, they primarily yield shorter-lived fission products such as cesium-137, with a half-life of 30 years. A process that could completely fission these long-lived TRU isotopes would greatly reduce the time the remaining waste would need to be isolated from the environment—but not enough to obviate the need for a geologic repository. For example, although cesium-137 would remain dangerous for only 300 years, instead of the 240,000 years needed for plutonium-239, geologic disposal would still be necessary, since one cannot assume that current institutions will remain viable and able to safely manage an interim surface storage facility for even that period of time.

But the long-lived TRU is only part of the problem. LWR spent fuel also contains long-lived fission products, such as iodine-129 (half-life: 15.7 million years), and technetium-99 (half-life: 211,000 years) that cannot be fissioned. For decades, elaborate schemes have been devised to attempt to separate such fission products and transmute them to stable isotopes, but none has been implemented. Even if ultimately successful, the cost and difficulty would be formidable (Chiba et al. 2017). These fission products would also need to be geologically isolated in a deep underground repository for as long as some TRU isotopes.

And in any event, it would not be practical for any real-world system to effectively reduce the entire inventory of TRU to the extent necessary to eliminate or even greatly diminish the need for long-term deep geologic repositories, contrary to the American Nuclear Society and Argonne

National Laboratory claims cited above. Notably, it is impossible to eliminate all the TRU in spent fuel—some fraction will inevitably end up in the waste stream and will require hundreds of thousands of years of geologic isolation. Nevertheless, if a system could reduce the quantity of TRU by a significant fraction—say, by 99 percent or more—this might enable a repository to meet less stringent safety criteria, reducing the cost and increasing the number of technically suitable sites. The process would also reduce the decay heat of the remaining waste in the long term. Depending on repository characteristics, this long-term heat reduction could potentially allow waste to be packed more densely in a repository, reducing the disposal space required per unit of electricity generated. (For the Yucca Mountain repository, in order to realize this benefit, the shorter-half-life elements cesium-137 and strontium-90 would also have to be extracted from the waste and stored above ground for 300 years—a questionable assumption, as discussed above.)

But if the amount of TRU that is ultimately left over is too large, then the benefits for repository disposal would not be great enough to justify the cost and security risks of reprocessing and recycling TRU. As discussed below, although the amount of TRU lost to waste streams is a critical factor, one also must consider the total TRU amount remaining in the system—including the reactor cores, fuel cycle facilities, and storage sites. If the system shuts down in the future, all of the remaining material would also need to be disposed of in a repository. But as shown below, the system would need to operate for hundreds or even thousands of years to reduce the total TRU inventory significantly. The present generation cannot guarantee that future generations will continue to operate, repair, and replace these systems for the length of time needed to achieve the necessary TRU reduction goal. If a reactor technology cannot significantly reduce the total TRU inventory in the system within a generation or two, future generations would still be stuck with a large stockpile of TRU—a situation only slightly better than the one that exists now.

INTERGENERATIONAL EQUITY

In addition to being impractical, a nuclear waste management strategy obligating future generations to maintain and operate a TRU burning system is inconsistent with the “intergenerational equity” principle. According to this principle, “those who generate the wastes should take responsibility, and provide the resources, for the management of these materials in a way which will not impose undue burdens on future generations,” and “a waste management strategy should not be based on a presumption of a stable societal structure for the indefinite future, nor of technological advance; rather it should

aim at bequeathing a passively safe situation which places no reliance on active institutional controls” (NEA 1995).

A robust geologic repository capable of containing waste for tens of thousands of years without the need for active controls and monitoring (beyond a reasonable period of retrievability) is arguably consistent with intergenerational equity. But a system requiring hundreds or thousands of years of costly and complex human activities to achieve its goals is clearly not. Our generation would bequeath to the future the obligation of maintaining and operating the system, without regard to cost and risk burdens. A TRU-burning system could only be consistent with intergenerational equity if it achieved its waste reduction goals within a few generations. The analyses discussed below show that even 120 years would not be sufficient.

SPENT FUEL “BURNING” REQUIRES REPROCESSING

It is also critical to realize that the term “waste burning” is an oversimplification that fails to convey the difficulty, cost, and risks of the industrial processes needed to extract re-usable materials from spent fuel and fabricate them into fresh fuel (see Box 3, p. 37).

BURNING THE TRANSURANIC ELEMENTS (TRU) IN NUCLEAR WASTE

Although complete destruction of radioactive waste is not possible, a key question is whether the TRU in nuclear waste can be reduced deeply and rapidly enough to significantly reduce the need for deep underground repositories. Typically, one of the limiting factors in a geologic repository is the heat load of high-level waste, and the precise limits for a given repository will depend on its geochemical characteristics and design. The TRU is the primary heat source in the waste after several hundred years, so a reduction in the TRU content of high-level waste is a necessary (but not sufficient) condition to pack more waste in a given repository volume. Another limiting factor is the long-term environmental contamination that will occur when a repository starts leaking radioactive material far in the future (because any repository will eventually leak over time). If more waste is packed into the same repository space, dose rates would increase and potentially exceed regulatory limits for public exposure, depending on the nature of the repository and many other factors.

To address this question, one must define what constitutes a “significant” reduction in TRU by a waste-burning system relative to LWRs operating on a once-through cycle. Analysts have used different standards over the decades, ranging from a reduction in total TRU mass by a factor of more than 1000 to as low as a factor of 10 (see Box 4, p. 38).

Unfortunately, even using the least stringent reduction factor of 10, which was adopted by the DOE in a 2009 study, it would take a very long time for TRU-burning systems in practice to have a meaningful impact on repository requirements (see Box 4, p. 38). This general result has been confirmed by many studies of fuel cycle systems, including a seminal National Academy of Sciences study (NAS 1996). In the appendix to this report, simple models are provided to illustrate this important finding.

WHAT LEVEL OF WASTE REDUCTION IS POSSIBLE?

In any real-world spent fuel reprocessing and recycling system, there are two primary sources of TRU-containing radioactive waste. First, there are process losses. Every time spent fuel from a reactor is reprocessed and refabricated into new fuel, a certain quantity of plutonium and other TRU end up in the waste streams. One can reduce that amount to very low levels, but that increases cost. Over time the mass of TRU that end up in unrecoverable waste streams can become significant, even if the waste from any one cycle is very small.

The second source is TRU within the system that remains unfissioned for practical reasons. Many analyses of the waste reduction benefits of reprocessing and recycling only account for the material entering and leaving the system; they ignore the nuclear material within the system. This is a huge oversight. There will always be TRU within a nuclear power system at any one time—in reactor cores, fuel fabrication plants, reprocessing plants, and interim storage facilities. However, unless one assumes that future generations will continue to operate the system (and replace old facilities) forever, eventually it will have to shut down, rendering unused fuel materials as radioactive wastes requiring geologic disposal. These materials need to be counted when estimating the overall reduction in TRU that the system can achieve. This observation was a key insight of the National Academy of Sciences study on separations and transmutation of nuclear wastes (NAS 1996).

The National Academy of Sciences evaluated the TRU reduction performance of a number of burner reactors and associated closed fuel cycles. The study found that if all the TRU in both wastes and operating facilities are considered, these systems will have to operate for an impractically long time—centuries or even millennia—to achieve a 100-fold reduction in the total mass of TRU.

These results have been confirmed by many other detailed systems analyses. A 2009 study by the Electric Power Research Institute and Electricité de France assessed the impact of phasing in a fast reactor system operating together with LWRs (35 percent fast reactors and 65 percent LWRs), while keeping the total US nuclear generating capacity

continued on p. 38

BOX 3.

Reprocessing and Recycling: Turning Spent Fuel into Fresh Fuel

By definition, spent fuel is discharged from a nuclear reactor when it can no longer be used as fuel in its current form. There are several reasons a nuclear fuel rod has a limited lifetime when irradiated in a reactor. The concentration of fissile material in the fuel decreases while the quantity of neutron-absorbing fission products increases—to the point when the fuel rod is unable to sustain a nuclear chain reaction. Irradiation increases the pressure inside the fuel rods (due to the generation of gaseous fission products) and decreases the strength of the metallic cladding surrounding the fuel rods until the rods are at high risk of rupture. A point is reached when it is no longer safe or productive to continue to irradiate the fuel rod.

What would it entail to use the spent fuel from LWRs as fuel for a new reactor? No reactor concept today can safely use spent fuel directly as new fuel. Instead, the spent fuel would have to undergo some type of processing—referred to as reprocessing—before it can be turned into fresh fuel.

First, mechanical and chemical processing would be needed to separate to some degree the fissile components of the fuel, such as plutonium and other actinides, from other spent fuel constituents. This can be an aqueous process, in which the spent fuel is dissolved in an acidic solution, or a non-aqueous process, such as conversion to metal (reduction) and electrometallurgical treatment (pyroprocessing). As discussed in chapter 4, reprocessing is costly, requiring the use of shielded facilities and remote handling equipment. It is environmentally hazardous. And most importantly, it increases the risk of nuclear proliferation and makes bomb-usable material easier for terrorists to steal.

After the spent fuel has been reprocessed to remove all the unusable or problematic isotopes, fresh fuel will have to be fabricated, with the new fuel form determined by the requirements of the type of reactor that will use it. Whether this process entails the fabrication of solid fuel or liquid fuel, it provides opportunities for diversion or theft of weapon-usable nuclear materials and therefore must be subject to stringent safeguards and security.

What Level of Transuranic Reduction in Radioactive Waste Would Make a Real Difference?

A comprehensive 1996 report by the National Academy of Sciences concluded that in order to “have a significant effect” on the total mass of TRU that would require geologic disposal, “an entire system of many facilities would be needed in which all the components operate with high reliability in a synchronized fashion for many decades or centuries. . . . The magnitude of the concerted effort and the institutional complexity . . . are comparable to large military initiatives that endure for much shorter periods than would be required” (NAS 1996). This report estimated in 1996 that the cost of such a system would be at least \$500 billion (or more than \$800 billion in 2020 dollars).

How great a reduction in the TRU inventory could justify the substantial expense of building and operating such a system over the decades or centuries that would be required? The National Academy of Sciences study pointed out that performance standards “changed markedly in recent years and have not been clearly defined,” but that the expectation of the DOE Advanced Liquid Metal Reactor Program in the early 1990s was a “thousand-fold reduction in the quantity of actinide waste going to a geologic repository” (NAS 1996). The study itself did not adopt a specific performance standard but implied that the authoring committee considered a 100-fold reduction as “significant” (NAS 1996).

In 2005, the DOE adopted the objective of achieving a 100-fold reduction in the quantity of TRU requiring disposal as one of the programmatic goals for its Advanced Fuel Cycle Initiative, which sought to develop a spent fuel reprocessing and recycling infrastructure in the United States (Piet et al. 2011). The DOE observed that such a reduction would delay the need for a second geologic repository until the end of the 21st century while allowing for significant nuclear power growth.

In a more recent evaluation of fuel cycle options, DOE scientists used a less stringent criterion—an order of magnitude (factor of 10) or more—to define “significant improvements” in nuclear waste management and other nuclear power metrics relative to the LWR once-through cycle, including repository decay heat load (Wigeland et al. 2009). Their logic was that “significant benefits” should be those “resulting in an improvement that is clearly larger than the uncertainties . . . typically an order of magnitude or greater” (Wigeland et al. 2010).

While it is difficult to define an objective standard because these assessments are so complex and uncertain, the present report will reference the DOE factor of 10 as the standard for a significant reduction in TRU. However, this standard is questionable, given that the analyses used to calculate the actual TRU reduction in a given system have large uncertainties and are highly sensitive to various assumptions. For example, one study finds that the estimated increase in Yucca Mountain repository capacity gained from reprocessing and TRU recycling (which is a function of the TRU reduction factor) would decrease by a factor of 50 as the assumed separation efficiency of TRU and fission products from waste decreases from 99.99 percent to a more realistic 99.0 percent (Wigeland et al. 2006). Thus, it is not apparent that a calculated improvement of a factor of 10 would be “clearly larger than the uncertainties.”

Moreover, it is not clear that a TRU reduction factor of 100 or even 1000 would be sufficient to meet waste disposal objectives. The original goal of the DOE’s Advanced Liquid Metal Reactor Program was to reduce the quantity of TRU elements in a repository to below the release limits stipulated by 40 CFR 191, the Environmental Protection Agency rule for geologic repositories (other than Yucca Mountain), which would require a reduction in plutonium-239 by a factor of more than 3000 (NAS 1996). In the realm of hazardous waste disposal, the standard for effective destruction of toxics is even higher. For example, a factor of one million (a 99.9999 percent reduction) is used by the Environmental Protection Agency as a standard for the destruction of dioxin, a long-lived and highly hazardous substance. No TRU reduction scheme is capable of achieving such a dramatic goal, no matter how long the system is operated.

Lastly, transmutation is not the only means of achieving an increase in the capacity of a repository. For example, it has been estimated that the long-term decay heat reduction from a 1000-fold reduction in the TRU mass in high-level waste would allow only a five-fold increase in Yucca Mountain capacity unless cesium-137 and strontium-90 are removed for above-ground storage (Wigeland et al. 2006). Given that the physical capacity of Yucca Mountain may be as great as nine times the current legal capacity (and assuming that the project is still viable), changing the law would be far cheaper than developing a TRU-burning system but would have an immediate impact on the quantity of waste that could be buried there.

continued from page 116

constant (Machiels, Massara, and Garzenne 2009). The study found that the system would have to operate for 70 years to reduce the total TRU mass in the system by just a factor of two relative to the once-through cycle used by LWRs. To reduce the TRU mass by a factor of 10 would require continuous operation for 632 years.

It is clear from these estimates why a 2009 DOE study concluded that (assuming a factor-of-10 standard for significance) “continuous recycle appears to be the only practical fuel cycle strategy that can significantly affect waste management issues for [used nuclear fuel] and [high-level waste], but only if all of the TRU is recycled, leaving only fission products and residual amounts of TRU in the [high-level waste]” (Wigeland et al. 2009). Any leftover spent fuel would count against the system’s overall capability for TRU reduction. Or, as an Idaho National Laboratory article in 2011 simply put it, “significant material accumulates throughout the system during recycling; thus achievement of high waste management benefits depends on continuation of recycling. Do not stop!” (Piet et al. 2011). Therefore, the system would have to operate forever. Such an exhortation is clearly inconsistent with the intergenerational equity principle.

The basis for these conclusions can be illustrated through relatively simple models. In the appendix to this report, examples are provided that demonstrate why it takes so long for a fast reactor system to burn up a significant fraction of its TRU fuel. Using optimistic performance assumptions, a fast neutron reactor operating as a TRU burner for 120 years would only reduce the total amount of TRU in the system by a factor of around eight—below the DOE’s factor-of-10 standard for a significant reduction.

In summary, while the idea of burning nuclear waste sounds appealing on the surface, such burning cannot be done quickly or efficiently enough to be an effective waste management strategy. The marginal benefits of developing and deploying systems for TRU burning do not justify taking on the proliferation, security, and safety risks of reprocessing and recycling spent fuel.

Uranium Utilization Efficiency

LWRs that are operated on a once-through cycle use only about 0.6 percent of the uranium mined for their fuel for energy. The remainder—more than 99 percent—is contained in the reactor’s spent fuel (around 10 percent) and the depleted uranium produced by enriching natural uranium for reactor fuel (just under 90 percent).

Some argue that the LWR once-through cycle is inefficient and wasteful, and should be replaced by fast breeder

reactors and a closed fuel cycle with reprocessing (Lynas 2011). They claim that this unsustainable use of uranium will eventually deplete the resource.

This argument has been cited by the International Atomic Energy Agency. In June 2018, then-director general of the agency, the late Yukiya Amano, said that although “identified uranium resources are sufficient for well over 100 years of supply . . . the current over-supply may not last forever. It is therefore important that this vital resource is mined, produced, and managed sustainably.” He pointed to “promising work . . . underway on new generations of nuclear power reactors that require less uranium” (Amano 2018). Along the same lines, in a 2014 study, DOE researchers adopted the following objective for improved natural resource utilization in nuclear fuel cycles: “on a per unit energy basis, [a] reduction in the amount of fuel resources needed by a factor of 100 or more” compared to the once-through LWR fuel cycle (Wigeland et al. 2014).

But is it really critical for the future of nuclear power to develop fuel cycles that use less uranium? The cost of uranium is only a small component of the cost of electricity to begin with (WNA 2020). And the world is not in danger of running out of uranium any time soon, even if nuclear power expands according to the most recent IAEA “high case” projections, and uranium remains so cheap that there is no economic incentive to use it more efficiently. The up-front capital investment needed to build a fast breeder reactor system and associated fuel cycle facilities would be substantial, but significant benefits to a nuclear utility’s bottom line would not be realized until the price of uranium is far higher than it is today, which is likely to be a long time from now (see chapter 5).

Nevertheless, as discussed above, conserving natural resources could be a worthwhile goal even if not warranted by current market conditions. And there are other benefits to using uranium more efficiently. Doing so would reduce the need for uranium mining, which is dangerous for workers and pollutes the environment, and would perhaps even reduce the demand for enrichment, a proliferation-sensitive part of the nuclear fuel cycle. However, the safety and proliferation risks of uranium mining and enrichment could also be reduced by more stringent regulatory controls and nuclear material safeguards—which would likely be less costly than developing and deploying more uranium-efficient reactors.

Thus, developing advanced reactors and fuel cycles that use uranium more efficiently is not essential for nuclear power’s future, but could be beneficial, provided they do not increase proliferation, terrorism, and safety risks and are cost-effective compared to alternatives for reducing the impacts of uranium mining.

IMPROVING URANIUM UTILIZATION

The uranium utilization efficiency of a reactor is generally defined as the ratio of the amount of heavy metal (e.g., uranium or plutonium) that undergoes fission (and hence releases energy) over the reactor's lifetime to the mass of natural uranium that was used to produce all the reactor's fuel. The amount of heavy metal fissioned includes both the direct fission of uranium isotopes (primarily U-235) and the fission of plutonium and other TRU produced through neutron absorption by U-238 and heavier nuclei. Since there is a very low probability that U-238, the major component of natural uranium, will undergo fission, to use natural uranium most efficiently, a reactor and fuel cycle system must convert as much U-238 to plutonium as possible and then fission as much of that plutonium as possible to release energy.

To calculate uranium utilization efficiency correctly, it is necessary to account for the entire amount of natural uranium used to produce all the fuel a reactor will need over its lifetime. This includes not only fuel that is periodically fed into the reactor when it reaches steady-state operation, but also the fuel for the startup core and for the intermediate cycles during the transition to steady-state operation. The latter contribution is particularly important for some reactor systems that can take many years or even many decades to reach a steady state. The uranium utilization efficiency is then defined as the total amount of heavy metal fissioned over the lifetime of the reactor divided by the total amount of natural uranium required (Yu et al. 2019). Also if a reactor uses thorium fuel, as do some NLWRs discussed in the present report, then the amount of natural thorium required should also be included. In that case, the parameter of interest is the "natural resource" utilization.

There are two reasons why the uranium utilization efficiency of current LWRs is so low. First, as discussed in chapter 2, the amount of heat energy that can be extracted from a given mass of fuel (the burnup) is limited. Second, the enrichment process results in the generation of a large stockpile of depleted uranium that is not usable as LWR fuel (and is thus a waste product) unless it is enriched, which is not economical as long as the uranium price remains low.

One reason why the burnup of fuel in LWRs is limited is that the proportion of fissile isotopes decreases as the fuel is irradiated. An insufficient quantity of new fissile isotopes, such as plutonium-239, are generated to compensate for the reduction in the quantity of U-235 in the fresh fuel that is fissioned.

As discussed above, spent fuel discharged from an LWR at a typical burnup of around 50,000 megawatt-days of thermal energy per ton of heavy metal has a U-235 content of less than

1 percent by weight and a total plutonium content of just over 1 percent by weight. The spent fuel also contains about 0.1 percent of other TRU, such as americium-241. The balance, around 93 percent, is almost entirely U-238. This means that 95 percent of the remaining heavy metal in the fuel (primarily U-238, U-235, and plutonium) was not used to produce energy and is contained in the waste. In addition, the leftover U-238 in spent fuel is only a fraction of the unused U-238 in the LWR fuel cycle. When natural uranium is enriched in the U-235 isotope for producing LWR fuel, a large stockpile of depleted uranium (containing greater than 99.7 percent U-238) is created, which is typically discarded as waste. A typical 1000 MWe LWR operating at 90 percent capacity and an 18-month refueling cycle requires around 20 metric tons of LEU fuel each year (at 4.5 percent U-235). About 180 metric tons of natural uranium would be enriched to produce this fuel annually. This reactor would fission a little more than one metric ton of heavy metal per year. The uranium utilization is therefore about 1 metric ton/180 metric tons = 0.6 percent.

However, this refueling strategy, which is typical for US LWRs today, is not optimized for efficient uranium use, but instead for maximizing the capacity factor by increasing burnup. Higher burnup fuel can be used for a longer time, increasing the cycle length and decreasing the average outage time for refueling between cycles. The same LWR in the previous example could operate on a yearly refueling cycle with a smaller required uranium enrichment (3.3 percent U-235) and fuel burnup. This refueling strategy would use only about 160 metric tons of natural uranium feed to generate the same amount of energy, increasing the uranium utilization by about 10 percent.

This example illustrates an important fact: uranium utilization is not necessarily improved if burnup is increased only through using higher enrichment fuels, because more depleted uranium will also be generated. This is why high-temperature gas-cooled reactors do not use uranium more efficiently than LWRs even though their fuel can achieve higher burnups than LWR fuel (see chapter 6). This is also true for conventional sodium-cooled fast reactors such as the TerraPower Sodium (see chapter 5).

How then can a fast reactor extract 100 times the amount of energy from a given quantity of uranium ore as an LWR does, as Argonne National Laboratory claims? This is only possible for a fast reactor operating in a breeding mode in a closed fuel cycle with reprocessing. Recall that the quantity of uranium and plutonium in LWR spent fuel comprises only about 10 percent of the mass of uranium mined to produce the fuel, while the remaining 90 percent is primarily bound up in depleted uranium tails. To achieve a 100-fold

increase in efficiency—that is, a uranium utilization of 60 percent—the reactors not only would have to fission all of the uranium and plutonium in LWR spent fuel, but also would have to convert 50 percent of the depleted uranium generated to produce the LWR fuel into plutonium and fission it completely. As shown below, this is a formidable task in practice.

THE FAST BREEDER FUEL CYCLE

The goal of the breeder reactor fuel cycle is to maximize natural uranium utilization by converting as much U-238 to plutonium as possible. However, to fully utilize the U-238 in LWR spent fuel in a fast reactor, it must be separated from the spent fuel through reprocessing, fashioned into targets (known as blankets), loaded into the reactor, and bombarded with neutrons. The depleted uranium tails also must be processed and fabricated into blanket fuel, but do not require reprocessing because the material is not irradiated.

Since the blanket material alone cannot sustain a chain reaction, the fast reactor must also be loaded with driver fuel. The preferred fissile fuel for a fast breeder reactor is plutonium, which can be obtained from reprocessing LWR spent fuel. Plutonium fission produces more extra neutrons in a fast spectrum that can be used to convert U-238 to additional plutonium. It is theoretically possible to breed more plutonium (and other TRU) in a fast reactor cycle than are consumed through fission (hence the name “breeder reactor”). The excess TRU generated by a breeder could be used as startup fuel for a new reactor.

The plutonium and other TRU bred in the blankets, as well as leftover uranium, would then be separated by reprocessing and used to fabricate fresh fuel. More blankets would then be loaded into the core and the process repeated. However, the process would not become self-sustaining until the system had reached a steady state, which could take several operating cycles.

Before a fast breeder reactor system becomes self-sustaining, it would need an external supply of plutonium obtained by reprocessing LWR spent fuel. But the process of enriching the fresh fuel needed for the LWRs generates a huge stockpile of depleted uranium. This stockpile must be accounted for in assessing the true uranium utilization efficiency of the system.

URANIUM UTILIZATION EFFICIENCY OF THE FAST BREEDER SYSTEM

In theory, the fast breeder system could achieve 100 percent uranium utilization efficiency, but only if the spent driver and blanket fuel from the fast reactor were repeatedly reprocessed; all of the recovered uranium, plutonium, and other TRU was recycled; and all of the U-238 contained in the original ore were converted to fissionable material that is fissioned to produce energy. (In practice, as with burner reactor cycles,

the unavoidable process losses—TRU that is discharged to the waste stream without being converted to energy—provide an upper bound to the utilization efficiency.)

However, using the model given in the appendix, one can show that it would take a very long time for a breeder reactor to convert a large fraction of the initial stockpile of depleted uranium to plutonium and utilize it to generate energy. Since the reactor would require only a small amount of the depleted uranium stockpile each year, the system of reactors and fuel cycle facilities could only utilize the entire quantity if they were rebuilt periodically and operated continuously for thousands of years.

According to the model (see appendix), 14,750 metric tons of natural uranium would have to be mined and enriched to fuel the LWRs needed to produce the plutonium for the initial cores of a 1000 MWe fast breeder reactor. At steady-state, the system would only require an input of 1.1 metric tons of depleted uranium each year. At this rate, the depleted uranium stockpile could fuel this fast reactor for nearly 12,000 years. At first glance, this seems like an amazing resource. And since the 1000 MWe reactor (operating at an 85 percent capacity factor) would fission about 0.8 metric ton of heavy metal per year, it would appear that the reactor’s uranium utilization efficiency is about 80 percent (0.8 metric tons of fission/1.1 metric tons of uranium).

However, it should be clear now that this is a misleading picture. According to the definition of uranium utilization efficiency presented above (Yu et al. 2019), the total amount of mined uranium used to produce the plutonium fuel for the fast breeder reactor over its lifetime must also be included. Using this definition, uranium utilization efficiency of the 1000 MWe PRISM fast breeder reactor operated for a 60-year period at 85 percent capacity would be about 50 metric tons of fissioned heavy metal/14,750 metric tons of natural uranium, or 0.34 percent—even less than that of an LWR. The annual uranium utilization would be less than 0.006 percent. For a breeder reactor to achieve a uranium utilization efficiency of 60 percent as claimed by Argonne—a 100-fold increase in efficiency of uranium use over LWRs—future generations would have to continue to operate, maintain, and replace fast breeder reactors and their associated fuel cycle facilities for thousands of years (around 11,000 years in the above example).

This is similar to the example for a TRU burner fuel cycle discussed earlier, which would also require hundreds or thousands of years to achieve its performance goal. But if a future generation were to decide not to continue to building and operating breeder reactors, then the remaining depleted uranium stockpile would not be utilized. Instead of a resource, it once again would be rendered a waste product.

Analyses of more detailed models confirm that the actual uranium utilization rate during a transition to a fast-breeder based reactor system would be far below 100 percent for hundreds of years. Department of Energy researchers have shown that when dynamic considerations are taken into account, such as the time lag required for spent fuel to be cooled, reprocessed, and refabricated into fresh fuel, the rate at which fast reactors can replace LWRs is significantly lower than indicated through static calculations (Piet et al. 2011). Consequently, even for the highest breeding ratio fast reactor considered (BR=1.75), the uranium utilization of the system would be at best no more than twice that of the LWR once-through cycle (1.2 percent) by the year 2100 and only 10 times more (6 percent) by 2200 (Piet et al. 2013). This assumed, optimistically, that all fast-reactor fuel would be reprocessed and recycled on site within a two-year period. For a more realistic 11-year lag period, the analysis found that the uranium utilization efficiency would only be 1.3 and 1.9 times more than the once-through cycle by 2100 and 2200, respectively. When process losses of uranium and other actinides to waste from reprocessing and fuel fabrication are taken into account, the closed fuel cycle becomes even less uranium-efficient.

BURNERS ARE NOT GOOD BREEDERS, AND VICE VERSA: THE DIFFICULTY OF MEETING BOTH SUSTAINABILITY GOALS SIMULTANEOUSLY

As discussed earlier, some fast reactor developers (such as Argonne National Laboratory) claim that their systems can simultaneously achieve two main goals of sustainability: the ability to greatly increase uranium utilization and to recycle the TRU in spent fuel. In the appendix to this report, it is shown why that is not possible. A system can be optimized either for increased uranium utilization (breeding) or for recycling TRU (burning), but cannot do both effectively at the same time.

The reason for this is simple: to use uranium most efficiently, a reactor and fuel cycle system must convert as much U-238 to plutonium as possible and then fission as much of that plutonium as possible to release energy. In contrast, to most efficiently fission TRU extracted from LWR spent fuel, the system must convert as *little* U-238 into new plutonium and other TRU as possible. TRU burners use only slightly (about 25 percent) less uranium than LWRs (Piet, Hoffman, and Bays 2010). The present study was unable to identify a system that could meet the criteria for significant reductions in TRU mass and efficient use of natural uranium simultaneously.

The reactors and fuel cycles described above represent attempts to improve nuclear power sustainability through reprocessing and recycling spent fuel. However, as discussed in chapter 4, those activities raise serious nuclear

proliferation and terrorism concerns, by rendering weapon-usable materials susceptible to theft or diversion. In light of that, it is prudent to consider ways to improve the sustainability of the once-through cycle without increasing proliferation and terrorism risks. If the current once-through fuel cycle could be modified to be more sustainable, the most compelling arguments for adopting a closed fuel cycle would no longer apply.

As discussed in chapter 8, it may be possible to develop reactors that can use uranium much more efficiently than current reactors on a once-through basis. Unfortunately, it is not likely that this can be accomplished by improving LWR technology because of the physical limits on burnup for conventional uranium oxide fuel and cladding. But as seen above, even if an LWR could fission 100 percent of its fuel, that would amount to only about 10 percent of the amount of natural uranium mined to produce the fuel. The remainder would be the depleted uranium left over from the enrichment process, which cannot be used in LWRs without re-enriching it or adding other fissile materials obtained from reprocessing, such as plutonium. Moreover, increasing fuel burnup alone does not increase uranium utilization if higher levels of enrichment are needed to enable higher burnup, because then even more depleted uranium would be generated (Kim and Taiwo 2010).

Therefore, to substantially increase uranium utilization in the once-through cycle, reactor systems would have to increase fuel burnup without an increase in the required level of uranium enrichment. Compared to current LWRs, such reactors would need to convert more U-238 to plutonium and fission more of that plutonium for energy. In addition, the reactors would have to be able to use U-238 as a fuel material, in order to utilize the inventory of depleted uranium tails. This approach is referred to as breed-and-burn. To date, the only reactor designs shown in theory to be capable of true breed-and-burn operation are fast reactors because extra neutrons are available for U-238 conversion in a fast spectrum. The TerraPower traveling wave reactor, which is a liquid sodium-cooled fast reactor, is the most prominent example. The TerraPower Sodium reactor, which is a once-through fast reactor with a conventional refueling cycle, is less uranium-efficient than an LWR (see chapter 5).

As discussed further in chapter 8, the uranium utilization efficiency of a successful breed-and-burn system would compare favorably to that of a fast-breeder fuel cycle, but without the need to separate and recycle weapon-usable TRU. In our assessment, the avoidance of reprocessing is a major selling point for breed-and-burn reactors. However, significant technical and safety challenges remain, and it is not clear at this time whether such reactors will be viable.

Nuclear Proliferation and Terrorism

Risks of Nuclear Power

Technologies for generating peaceful nuclear power are dual-use: they can also be used to produce the materials needed to make nuclear weapons. Nations that possess civilian uranium enrichment or spent fuel reprocessing technologies also have the means to produce fissile materials for nuclear weapons. The landmark Nuclear Non-Proliferation Treaty allows the 186 non-nuclear weapon state parties to possess dual-use nuclear facilities for peaceful purposes, but prohibits those states from acquiring nuclear weapons. The International Atomic Energy Agency (IAEA) is tasked with implementing a safeguards system to verify that nations are not diverting nuclear materials from declared nuclear facilities which could be used to produce nuclear weapons. In some states, the IAEA also has the authority to verify the absence of undeclared facilities or materials. The IAEA conducts safeguards inspections intended to detect the diversion of such materials in a timely manner.

The five nuclear weapon states—the United States, Russia, France, China, and the United Kingdom—are not obligated to accept IAEA safeguards, but can volunteer individual facilities for safeguards by placing them on an “eligible facilities list.” The IAEA, however, does not generally implement safeguards in nuclear weapon states due to a lack of resources.

The IAEA applies safeguards to “special fissionable materials,” which consist of enriched uranium, plutonium, and uranium-233 (U-233), as well as source materials such as natural uranium that can be used to produce special fissionable materials. Of these, it defines highly enriched uranium (HEU) and plutonium (containing less than 80 percent plutonium-238) and U-233 as “direct use materials.” HEU and plutonium are the materials that are believed to fuel all of

the world’s nuclear arsenals, but U-233 is also weapon-usable. In addition, certain transuranic isotopes such as neptunium-237 and americium-241 are weapon-usable, but they are considered “alternative” nuclear materials and are not in the scope of IAEA safeguards. Instead, the IAEA requests that states voluntarily track and report information about any stocks they possess.

As discussed in chapter 2, natural uranium consists primarily of a mixture of two uranium isotopes: U-238 (99.3 percent) and U-235 (0.7 percent). HEU is uranium with 20 percent or more of the isotope U-235; low-enriched uranium (LEU) has less than 20 percent U-235. It is extremely impractical, but not impossible, to make a nuclear weapon with LEU enriched to above about 10 percent U-235.

Current light-water reactors (LWRs) use LEU enriched up to about 5 percent U-235. However, facilities that enrich natural uranium to produce LEU for use in power plant fuel can be readily reconfigured to make HEU. Consequently, civilian enrichment facilities in Nuclear Non-Proliferation Treaty non-nuclear weapon states (and in some nuclear weapon states as well, on a voluntary basis) are under IAEA safeguards to verify that they are not being misused to produce HEU.

The potential for Iran to use its uranium enrichment facilities to develop nuclear weapons was a major reason for the international concern that led to the 2015 Joint Comprehensive Plan of Action—the so-called Iran Deal. While Iran had announced its intention to produce only LEU, it could be capable of using its enrichment facilities to produce HEU. However, the physical limits on its uranium holdings and enrichment capacity that were stipulated in the Joint Comprehensive Plan of Action increased the “breakout” time it would have taken Iran to produce enough HEU for a nuclear

weapon. After the United States withdrew from the agreement in 2018 and launched an airstrike that killed Iranian general Qassem Soleimani in January 2020, Iran announced that it would no longer abide by the Joint Comprehensive Plan of Action's operational restrictions on its nuclear program that had increased its breakout time (Zarif 2020). Although repairing the damaged relationship between the United States and Iran will be difficult, hopefully the Biden administration will be able to salvage the Iran Deal and restore its constraints on Iran's nuclear program.

A second route for nuclear proliferation is provided by civil reprocessing facilities, which use chemical processes to extract plutonium from spent reactor fuel. While the purpose of civil reprocessing is to separate plutonium that could be used in fresh nuclear reactor fuel (the "burning" of nuclear waste, as discussed in chapter 3), such plutonium could also be used to produce nuclear weapons. Indeed, nuclear reactors and reprocessing were first developed during the Manhattan Project to produce plutonium for nuclear weapons.

Reprocessing facilities present a greater risk than enrichment plants that produce only LEU. While a large uranium enrichment plant could be readily reconfigured to produce HEU, it would be difficult to do covertly, and there would be some time delay before sufficient HEU were available. In contrast, a nation that possesses a stockpile of separated plutonium has a readily available supply of material that it could immediately use to produce nuclear weapons, should it decide to do so.

Moreover as discussed below, it is feasible to surreptitiously divert enough plutonium to build a nuclear weapon from a commercial-scale reprocessing or plutonium fuel fabrication plant without timely detection by the IAEA. At such facilities, which could process several tons of plutonium each year, the measurement uncertainty alone could amount to far more than the relatively small amount of plutonium needed to produce a weapon.

In addition to its proliferation risk, reprocessing spent nuclear fuel makes it easier for terrorists to steal weapon-usable plutonium. In contrast to separated plutonium, un-reprocessed spent fuel contains highly radioactive fission products and must be shielded from human access and handled remotely using specialized equipment. For this reason, spent fuel is considered "self-protecting" in that anyone trying to steal it in a practical scenario would likely receive a high enough dose of radiation to cause a serious or even fatal injury. The Nuclear Regulatory Commission's current threshold for "self-protection" is 100 rem of radiation per hour for someone standing three feet away. The IAEA standard is similar but uses a distance of 1 meter, or 3.3 feet. Although the radioactivity of spent fuel declines over time, under

these standards spent fuel typically remains self-protecting for about 100 years.

An additional reason that spent fuel is difficult to steal is that it exists in the form of large and heavy assemblies. Today's pressurized-water reactors typically have fuel assemblies that are roughly 15 feet long and weigh 1000 pounds, or about 450 kilograms (kg). A single pressurized-water reactor fuel assembly contains around five kg of plutonium: about one percent by weight. It is the combination of its radioactivity, size, and dilution that makes spent fuel an unattractive target for thieves seeking to obtain plutonium.

Thus, reprocessing extracts plutonium from an object that is very difficult to steal and converts it to a form that could be more easily stolen while it is being processed, stored, or transported. Theft could be carried out by an external attacking force, an insider, or both working together. A covert insider theft might not be detectable for a long time because it is not possible to precisely measure the plutonium in a reprocessing facility, and therefore determine whether any is missing, while the plant is operating.

International Standards for Detecting Diversion of Nuclear Materials and Protecting Them from Theft

Because peaceful nuclear technologies can be misused, stringent controls are needed to help ensure that civil nuclear power does not facilitate the spread of nuclear weapons to non-nuclear states or terrorists. These controls include (1) international safeguards to detect and thereby deter nuclear proliferation by countries, and (2) domestic security and nuclear material accounting measures to deter theft by sub-national terrorist groups.

SAFEGUARDS AND MATERIAL ACCOUNTANCY

International safeguards are applied by the IAEA in countries with which it has legal agreements to verify that nuclear materials are not being diverted for undeclared use. This includes all the non-nuclear-weapon states that are party to the Nuclear Non-Proliferation Treaty that have nuclear facilities. In addition, the five nuclear weapon states (United Kingdom, China, France, Russia, and the United States) have voluntarily accepted safeguards on some of their civil facilities.

The IAEA's fundamental objective is "timely detection" of the diversion of a "significant quantity" (SQ) of weapons material—nominally the approximate quantity needed to make a first-generation nuclear weapon, taking into account process losses. For instance, it should be able to detect the abrupt diversion of one SQ of plutonium or HEU within one

month. The SQ value for plutonium is 8 kg; for HEU, it is a quantity of total uranium containing 25 kg of U-235.

The foundation of safeguards is material accountancy: the measurement of a facility's material inputs, outputs, and in-process inventory, to determine whether there is any "material unaccounted for" (MUF) and how large it is. Because all techniques to measure or estimate quantities of nuclear materials have uncertainties, and because nuclear material can get stuck in parts of facilities that are hard to access, there will always be a non-zero value of MUF. The challenge for IAEA safeguards inspectors is to determine, on a statistical basis, whether a given MUF represents an actual diversion of material or whether there is an innocent explanation for it.

Trying to detect a genuine diversion of one SQ of plutonium in a timely manner is a tough job at large commercial reprocessing plants, which separate many hundreds of SQs each year. For instance, at the still-unfinished Rokkasho reprocessing plant in Japan, which was designed to separate 8000 kg of plutonium each year, a diversion of more than 25 SQs—enough to make three first-generation nuclear weapons—would have to occur before the IAEA could conclude with 95 percent confidence that the resulting MUF was not due to a statistical error and therefore could be a sign of a diversion.

Over the last 25 years, several examples of large plutonium MUFs that went undetected for months or even years have come to light at plutonium-processing facilities around the world. These include the Tokai Reprocessing Plant in Japan in 2003 (206 kg of plutonium), the Thermal Oxide Reprocessing Plant (THORP) in the United Kingdom in 2005 (190 kg), and the Cadarache plutonium fuel production plant in France in 2002 (39 kg) (Kuperman, Socolow, and Lyman 2014). These examples underscore the inherent difficulty of achieving safeguards goals at such bulk-handling facilities.

Because the IAEA cannot meet its detection goals at bulk-handling facilities using material accountancy alone, it supplements this with "containment and surveillance" measures. These measures, which include closed-circuit television cameras and seals on nuclear material containers, are intended to ensure that no unauthorized movement of nuclear materials has taken place. However, these measures cannot fully compensate for inaccurate or slow material accountancy measures, as discussed below.

SECURITY

Unlike international safeguards, protecting nuclear facilities from sub-national terrorist attacks, such as theft of weapon-usable materials or radiological sabotage—is regarded by the

international community as a sovereign responsibility. Security measures include "guns, guards, and gates" to protect against external threats, as well as measures to mitigate insider threats such as background checks for personnel. Nuclear plant security is also increasingly being challenged by emerging threats such as cyberattacks and malevolent use of aircraft such as drones. States are also responsible for maintaining material accountancy measures so that nuclear facility operators can determine whether terrorists are stealing nuclear material, or to quickly resolve claims of theft that could be used for blackmail. (The accounting systems developed by states also play a dual role by providing data to IAEA inspectors for verifying compliance with their safeguards agreements, although state regulators may have different performance standards for those systems than the IAEA.)

The Nuclear Regulatory Commission (NRC), which has oversight over US commercial nuclear facilities, classifies the most sensitive nuclear materials containing plutonium, enriched uranium, and U-233 as Category I, II, and III, and has developed security standards for each category. These categories depend on the type of nuclear material, the quantity, and whether the material is irradiated to the self-protection standard defined above (but not on other factors such as whether the material is pure or diluted with another substance). The highest level of physical protection, Category I, is applied to certain quantities of materials that can be directly used in the manufacture of nuclear weapons. For example, 2 kg or more of unirradiated plutonium, and 5 kg or more of U-235 contained in HEU, fall under Category I. In contrast, the highest security category for low-enriched uranium with a U-235 content below 10 percent—which includes LWR fuel—is Category III. And 10 kg or more of uranium with a U-235 content from 10 percent to below 20 percent—defined in this report as high-assay LEU (HALEU)—is considered a Category II quantity, with an intermediate security risk.

The Department of Energy (DOE), in managing the nuclear materials under its control, uses a more complex security scheme that, in addition to considering the type and amount of material, also takes into account the material's physical and chemical properties. These characteristics are relevant to the material's attractiveness to someone seeking to build a nuclear bomb. Attractiveness is based on considerations such as whether the material could be used to make a weapon directly or would need further refinement. And if the material requires refinement, the attractiveness ranking accounts for how difficult and hazardous it would be to steal, transport, and process it. Such considerations, however, are subjective and depend on the assumed capabilities of terrorists.

The Nuclear Non-Proliferation Treaty and IAEA safeguards do not include prevention of nuclear terrorism within their scope, in accordance with the belief that security should be a national responsibility. A different international instrument, the amended Convention on the Physical Protection of Nuclear Material, obligates its parties to ensure that nuclear facilities within their borders, as well as international transports, meet a basic set of security standards. The convention incorporates a material security categorization scheme similar to that of the NRC. However, unlike the Nuclear Non-Proliferation Treaty, the convention does not contain any mechanisms for the enforcement of its provisions.

NLWRs: Enrichment Issues

Weapon-usable HEU has been commonly used as a fuel for NLWRs. In the past, fast reactors such as the Experimental Breeder Reactor-II (EBR-II) in Idaho; molten salt reactors such as the Molten Salt Reactor Experiment; and high-temperature gas-cooled reactors, including Fort St. Vrain, have all used HEU. Fast reactors operating today, including the Russian BN-600 and BN-800, continue to use full or partial cores of HEU fuel. Historically, designers preferred the use of HEU for technical reasons and discounted the security risks. Today, however, there is greater awareness of the proliferation and terrorism risks of HEU use. Deploying new reactors that use HEU would violate a growing international norm discouraging the civil use of HEU. As a result, NLWRs that might have used HEU in the past are being designed today to use LEU instead.

However, many NLWR designs today require high-assay LEU (HALEU) because they need higher fissile enrichments and burnups than LWRs (see chapter 2). Recall that HALEU is defined as LEU with an enrichment from 10 percent to below 20 percent U-235, in accordance with a working definition used by the URENCO uranium enrichment consortium (see endnote 8). Proposed reactors that would use HALEU include the Oklo liquid-metal-cooled fast micro-reactor, the Xe-100 pebble-bed high-temperature gas-cooled reactor, and the ThorCon thermal molten salt reactor, discussed in chapters 5, 6, and 7, respectively.

Although HALEU is LEU and is not considered practical for use in nuclear weapons without further enrichment, it does present additional nuclear proliferation and security risks compared to LEU with lower enrichments. A key issue is that current NLWR designs would require large quantities of HALEU. For example:

- A 1 gigawatt-electric (GWe) ThorCon plant, consisting of four reactor modules, would require 9.44 metric tons of 19.75 percent-enriched HALEU for the initial cores,

and a supply of 2.63 metric tons per year over the eight-year reactor lifetime. This corresponds to an average of 3.8 metric tons per GWe-year. (Jack Devanney, principal engineer of the ThorCon molten salt reactor, email message to the author, January 4, 2018.)

- Each Xe-100 76 megawatt-electric (MWe) module would require 1.5 metric tons of 15.5 percent-enriched HALEU for its first core and an annual supply of nearly 0.5 metric ton per year. This corresponds to a requirement of about 6 metric tons per GWe-year.
- Based on information provided in the Oklo “Aurora” 1.5 MWe license application to the NRC, the reactor would operate without refueling for 20 years and the peak fuel burnup would not exceed 1 percent (Oklo 2020). This indicates that the core would require at least 3 metric tons of HALEU, corresponding to a relatively high HALEU demand of 100 metric tons per GWe-year. This is consistent with the published requirements for a similar reactor concept, the Los Alamos Megapower reactor: 4.6 metric tons of 19.75-enriched HALEU for a five-year core lifetime, which works out to 460 metric tons per GWe-year.

Project Pele, the Department of Defense’s mobile micro-reactor program, could also require a significant supply of HALEU if it moves forward with prototype micro-reactor demonstration and deployment. The project specifies that the reactors must use HALEU fuel. While those very small reactors (10 megawatt-thermal or less) would likely have higher-burnup fuel than Oklo’s Aurora, they would still require substantial quantities of HALEU—likely many hundreds of kilograms over their operating lives.

Annual HALEU demand for a reasonably sized fleet of NLWRs such as ThorCon or the Xe-100 could easily be hundreds of times greater than the current rate of supply. The Nuclear Energy Institute recently projected that the US nuclear industry could need more than 200 metric tons of HALEU per year by 2031 (Redmond 2020). In contrast, demand for US-origin HALEU by foreign research reactors, produced by down-blending excess military HEU stocks with natural uranium, is only around 1.5 metric tons per year. Even at that low rate, the current supply of excess HEU that the United States has designated for converting to HALEU will be exhausted around 2040 (Lyman 2018a).

It is unlikely that sufficient additional military HEU would be available for downblending to HALEU for the first demonstration reactors, much less for a commercial fleet. Therefore, new uranium enrichment capacity for HALEU would be required, either domestically or internationally. In addition to enrichment facilities, a fleet of HALEU-

fueled NLWRs would require a new fuel cycle infrastructure including conversion, fabrication, waste management, and (eventually) disposal facilities.

This infrastructure is not likely to be limited to the United States. If the United States moves forward with commercialization of HALEU-fueled NLWRs, US companies will likely seek to export them, and the rest of the world may pursue their own programs. Brazil has already expressed interest in producing HALEU at its domestic uranium enrichment plant (Guimaraes and Perrotta 2020). The production, processing, and transport of large quantities of HALEU around the world could pose significant risks of nuclear proliferation and terrorism if not appropriately safeguarded and protected.

Nuclear Terrorism and Proliferation Concerns of HALEU

Although the direct use of HALEU in a nuclear weapon would be impractical, its production and use on a large scale could have significant implications for the risks of nuclear terrorism and nuclear proliferation (Lyman 2018a). These issues need to be thoroughly assessed before the United States goes forward with an NLWR development program that could stimulate a global demand for the material.

As discussed below, the enhanced security risk of certain quantities of LEU with an enrichment greater than or equal to 10 percent but below 20 percent (defined here as HALEU) is reflected in domestic and international material security standards. That factor alone would increase security costs for reactors that use HALEU and their associated fuel cycle facilities.

There are two main reasons why the nuclear terrorism and proliferation risks of HALEU are greater than those of lower-assay LEU. The first is that the material can be used directly in nuclear weapons. The second is that it is somewhat easier to enrich it to HEU.

The first reason cannot be addressed definitively here because there is very little public information on the effort needed to use HALEU directly in a nuclear weapon. However, it appears from available information that although it would likely be highly difficult for nations or terrorists with unsophisticated nuclear weapons programs, it is not considered impossible. The former director of the Theoretical Division at Los Alamos National Laboratory, J. Carson Mark, testified at a 1984 congressional hearing that “it is possible on paper to imagine that you could make an explosive out of anything in that [21 to 90 percent enrichment] range, and in fact, it’s even possible down to 10 percent” (Mark 1984). However, he went on to say that “the penalties are

quite tremendous.” For one, the total quantity of uranium that would be needed to make a bomb is considerably higher at lower enrichments: about 10 times higher at 20 percent than at 90 percent. But even so, the amount of 19.75 percent-enriched HALEU needed for a bomb could be around 300 kilograms. Thus, if a bomb with such a massive core were feasible, a single Oklo micro-reactor core would contain about 10 nuclear weapons’ worth of material.

More recently, a review by the DOE national laboratories of the attractiveness of various types of nuclear materials for use in nuclear weapons concluded that HALEU was of “low” attractiveness, defined as material that is “impractical, but not impossible” for a sub-national group to process and use in a nuclear explosive device (Ebbinghaus et al. 2013).

With regard to the second question, less enrichment effort would be required to produce weapon-usable HEU from HALEU feedstock than from the lower-assay LEU used in LWR fuel. For example, the production of 90 percent-enriched HEU would require about three times less separative work (a measure of the effort required to enrich uranium; see chapter 2) using 19.75 percent-enriched LEU feed than using 5 percent-enriched feed, and 1.7 times less than using 10 percent-enriched feed. Some analysts have argued that producing HEU from HALEU feed would require a relatively small enrichment plant that would be cheaper and could be easier to conceal than the plant needed to produce HEU from lower-enriched feed (Forsberg et al. 1998).

However, these differences in the amount of separative work needed to produce enough HEU for a weapon is not as significant for modern gas centrifuge plants, which are compact and scalable, as they may have been for older technologies such as gaseous diffusion plants. For countries with large commercial enrichment facilities producing LEU for LWRs, the availability of HALEU would not appear to make a big difference in the timeline for producing HEU if the country overtly violates its nonproliferation commitments, although it might be more beneficial for covert proliferation pathways, such as the use of a small clandestine facility to produce HEU from HALEU diverted from a declared facility. The advantages of access to HALEU would be greatest for a country with a relatively small enrichment capacity such as Iran.

PHYSICAL SECURITY REQUIREMENTS FOR HALEU

US domestic and international security regimes both consider HALEU to be a higher-risk material than lower-assay LEU. In the Convention on Physical Protection of Nuclear Material, the IAEA ranks HALEU as a more attractive material than lower-assay LEU for terrorists. Specifically, 10 kg or more of

U-235 is classified as Category II if contained in HALEU, but only as Category III if contained in low-assay LEU. This requirement dates to the earliest version of the IAEA's security recommendations in the mid-1970s, but it still applies today, despite advances in enrichment technology. The NRC, which adopted the IAEA's classification table decades ago, recently re-analyzed the relative attractiveness of different nuclear materials and reaffirmed the need to provide additional security for HALEU (Lyman 2018a).

Therefore, under current protocols, Category II security measures will be required for HALEU at nuclear fuel production facilities—measures more stringent than the Category III measures currently in place for LWR fuel facilities. The need for more robust security programs will have cost and management implications for NLWR reactors that use HALEU and the fuel cycle facilities that support them.

There are no licensed Category II fuel facilities in the United States. One challenge that will be encountered in the licensing of new Category II facilities and transport activities in the United States is the absence of updated security requirements for such facilities. In 2019, the NRC terminated a rulemaking that would have updated NRC security requirements for nuclear materials such as Category II facilities to address changes in the threat environment since the existing rules were promulgated decades ago. Consequently, Category II applications will have to be evaluated on a case-by-case basis, which could lead to inconsistent application of security upgrades. For that and other reasons, certain NRC staff formally objected to the NRC commissioners' decision to terminate the rulemaking.

SAFEGUARDS FOR HALEU

In contrast to the IAEA's framework for physical protection, IAEA safeguards do not distinguish between high- and low-assay LEU. Since HALEU is LEU, the IAEA considers it "indirect use material," and the SQ value and timeliness goal for detecting a diversion are the same as for lower-assay LEU: 75 kilograms of U-235 and 1 year, respectively. Nevertheless, the greater proliferation risks of stockpiling HALEU have been recognized by the international community. Iran committed under the November 2013 Joint Plan of Action to voluntarily reduce its inventory of "up to 20 percent"-enriched LEU to only what was needed for working stock for its research reactor and to temporarily not enrich above 5 percent—restrictions that were strengthened in the now-defunct July 2015 Joint Comprehensive Plan of Action but have now been violated by Iran. Arguably, these agreements created a new de facto safeguards category for HALEU that acknowledges that it does present a greater proliferation risk than lower-

assay LEU—but one that is not reflected in the IAEA's detection goals.

At first glance, it might appear that even under current guidelines the IAEA would need to apply more stringent material accountancy measures to meet its detection goals at a HALEU bulk processing facility than at a lower-assay LEU facility, because less material would have to be diverted to obtain 1 SQ. One SQ of LEU is 1.67 metric tons of total uranium for 4.5 percent-enriched LEU, but only 380 kilograms for 19.75 percent enrichment. Therefore, a HALEU facility might need a more sensitive safeguards system to detect the diversion of this smaller amount of material in a timely manner. However, the total amount of uranium corresponding to one SQ is a less important parameter than the fraction of facility throughput that it represents. If one compares fuel facilities sized to supply the same amount of nuclear power capacity per year, the detection requirements would be similar. A typical LWR fuel fabrication plant can supply about 1200 metric tons per year, or about 720 SQs: enough for about 60 1-GWe LWRs. In comparison, a HALEU fuel production plant supplying 60 GWe of X-Energy's Xe-100 reactors would have an annual throughput of about 360 metric tons of 15.5 percent-enriched HALEU, or about 740 SQs. Thus, 1 SQ would be about 0.14 percent of the annual throughput in either case, and diversion of 1 SQ over the course of a year without detection would be of comparable difficulty at both facilities. But, as discussed above, the consequences of a diversion of 1 SQ of HALEU, as currently defined, would be more serious than a diversion of 1 SQ of lower-assay LEU.

HALEU SAFEGUARDS AND SECURITY: WHAT IS NEEDED?

The fact that the IAEA recommends different security measures for LEU with enrichments below and above 10 percent, but treats all LEU as equivalent with regard to safeguards, is a troubling inconsistency that should be resolved if the production and use of HALEU expands. If HALEU requires more stringent security measures than lower-assay LEU, then it may warrant more intensive safeguards as well. US government agencies and the IAEA should take a hard look at the proliferation implications of a commercial HALEU fuel cycle and adjust their protocols accordingly.

In an ideal world, the IAEA would have the flexibility to introduce a smaller SQ (perhaps 50 kg) and a more stringent timeliness detection goal (perhaps six months) for HALEU to reflect its greater proliferation significance. Unfortunately, such radical changes are nearly impossible at the IAEA, given the reluctance of its international Board of Governors to approve more restrictive or intrusive safeguards obligations.

Short of that, the IAEA could consider separate, voluntary tracking of HALEU production and use, similar to its approach to the “alternative nuclear material” neptunium, at least until more stringent safeguards measures can be imposed.

Another issue to consider is the impact of NLWR deployment on global uranium enrichment requirements. Transitioning to a fuel cycle that requires less enrichment could be positive for nonproliferation by reducing the number and size of uranium enrichment plants needed around the world to support a given level of nuclear energy production.

For instance, once-through breed-and-burn reactors would, in theory, require less enrichment on average to support a given amount of electricity production. For example, TerraPower estimates that its traveling wave reactor would require, on average, only 25 percent of the uranium enrichment per unit of electricity required for current LWRs (Gilleland, Petroski, and Weaver 2016).

Unfortunately, other NLWR designs are not optimized for more efficient utilization of enrichment (Lyman 2018a). An LWR with a 60-year lifetime would require about 150,000 separative work units per GWe-year on average. In comparison, the X-Energy Xe-100 would require 190,000 separative work units per GWe-year on average, or 25 percent more than an LWR. And an Oklo-type fast micro-reactor would require 4 million separative work units per GWe-yr, or more than 25 times the LWR requirement—another indication why these reactors would not be practical for large-scale distributed power generation. Thus, deployment of some NLWRs that use HALEU would require an expansion of enrichment capacity to support a given level of nuclear electricity generation—a trend in the wrong direction.

Reprocessing and NLWRs

As discussed in chapter 3, spent nuclear fuel must undergo some type of chemical treatment, or reprocessing, before it can be used by any reactor designed to “burn” nuclear waste. Thus, any reactor concept that advertises an ability to burn spent fuel requires a fuel cycle that incorporates reprocessing.

Compared to the once-through LWR fuel cycle with direct disposal of spent fuel, all reprocessing technologies make weapon-usable materials such as plutonium much more vulnerable to diversion by countries or theft by terrorist groups seeking to obtain nuclear weapons. Fuel cycles with reprocessing require significantly greater resources than once-through cycles to pay for more intensive nuclear material accountancy, physical security, and (in non-nuclear weapon states) international safeguards activities. These additional

activities are costly because they require highly trained personnel and more specialized equipment.

Decades ago, in recognition of the dangers of reprocessing, the United States adopted a policy to not reprocess commercial spent fuel, with the goal of discouraging other nations from doing so. While this policy has shifted over the years (see Box 5, p. 50), the United States does not currently reprocess spent fuel from power reactors, and it has no firm plans to do so. But the DOE continues to fund research and development on reprocessing technologies and related advanced reactor projects, an indication that regardless of national policies and practices, there is strong support within the DOE and in Congress for developing a closed fuel cycle in order to recycle nuclear waste.

PYROPROCESSING

The standard reprocessing technology used worldwide today is PUREX. PUREX is an aqueous process, which begins with dissolving spent fuel in a water-based acidic solution. PUREX can be used for a variety of types of spent fuel materials, including oxide fuel from LWRs. However, certain fuels for NLWRs—such as the metallic fuel used in the GE-Hitachi PRISM fast reactor and fuels for molten salt reactors—are compatible with a different, non-aqueous type of reprocessing known as pyroprocessing.⁹ One often hears that pyroprocessing has lower proliferation and terrorism risks than PUREX. However, as discussed in detail below, this is a highly inaccurate claim.

Today, PUREX reprocessing takes place at a few centralized facilities, such as La Hague in France, which accept spent fuel from around the world. The plutonium separated from spent fuel is then shipped to fuel fabrication plants, for instance the MELOX plant, which is 700 miles away from La Hague. Fresh fuel containing plutonium is then shipped to reactor sites around the world. These transports are of particular concern because they are arguably the hardest to protect—and therefore most vulnerable—part of the nuclear fuel cycle.

In contrast, pyroprocessing facilities are far more compact than PUREX plants, making it feasible to incorporate on-site pyroprocessing and fuel fabrication plants into the reactor facility itself. The Integral Fast Reactor program (discussed in chapter 5) sought to develop a metal-fueled fast reactor with co-located pyroprocessing and fuel fabrication. Also, many molten salt reactor concepts would require co-located fuel reprocessing plants. Co-locating reprocessing and fuel fabrication plants at reactors would reduce the need to transport both spent fuel and fresh fuel, providing a security benefit compared to current reprocessing practices.

However, this benefit likely would be outweighed by the far greater risks presented by the large number of sensitive reprocessing and fuel fabrication facilities dispersed at multiple reactor sites. Nuclear reactor owners would be responsible for providing far higher levels of security than are needed for the reactors alone. Increased resources would also be needed by national regulators and international inspectors to safeguard distributed small reprocessing facilities. As discussed below and in chapter 7, molten salt reactors, which may require co-located pyroprocessing plants to

periodically treat the reactor fuel (or even continuously, depending on the design), would be particularly difficult to safeguard.

Therefore, if such reactor designs and their associated fuel cycle facilities were built in the United States, the risk of nuclear terrorism would increase by increasing the number of facilities possessing nuclear weapon-usable materials. Likewise, if they were built in other nations that do not have nuclear weapons, the risks of both nuclear proliferation and nuclear terrorism would increase.

BOX 5.

The Ups and Downs of US Reprocessing Policy

The United States began reprocessing spent fuel from US power reactors in the 1960s. This effort was reassessed after India's 1974 test of a nuclear weapon that used plutonium produced with reprocessing technology it had imported from the US under claims of "peaceful use." Under Presidents Ford and Carter, the United States adopted a no-reprocessing policy, arguing that the spread of commercial reprocessing facilities could spur the proliferation of nuclear weapons. They hoped the US policy would convince other countries to adopt a similar stance.

Several other countries, including Brazil, Pakistan, South Korea, and Taiwan, sought to follow India's example by launching ostensibly peaceful reprocessing programs. In each case, however, the effort was halted largely because of US opposition. The United States questioned these nations' motives for acquiring these technologies and argued that its own example showed that a robust nuclear power program does not require reprocessing.

The United States was also able to influence other countries through its agreements for nuclear cooperation, which allow it to export nuclear technologies and materials while retaining consent rights over the reprocessing of US-origin spent fuel. For example, the United States has to date blocked reprocessing in South Korea. (That decision will be subject to review in the future by a joint US-South Korean commission under the current nuclear cooperation agreement, which was renewed in 2015.) On the other hand, for political reasons, the United States has provided some other countries, most notably Japan, with blanket consent for reprocessing. This policy has resulted in Japan's accumulation of a stockpile of about 46 metric tons of plutonium—enough for thousands of nuclear weapons—which has caused concern around the world (Obayashi and Sheldrick 2018).

In 1981 President Reagan reversed the Carter administration's policy, allowing US companies to reprocess their spent fuel provided they paid for it, but the US industry did not do

so because it was too costly. In 1993 President Clinton reversed the policy once again, stating that "the United States does not encourage the civil use of plutonium and, accordingly, does not itself engage in plutonium reprocessing for either nuclear power or nuclear explosive purposes" (Clinton 1993).

In 2001, the George W. Bush administration's National Energy Policy called for a major expansion of nuclear power, along with a reconsideration of reprocessing and the use of plutonium for fuel. In 2006, the administration launched the Global Nuclear Energy Partnership, which would have entailed reprocessing US spent fuel and expanding the reprocessing capacity of certain partner nations. Under this plan, the United States and its partners would lease reactor fuel to other nations and require them to return the spent fuel for reprocessing, with the goal of dissuading them from acquiring their own enrichment and reprocessing facilities.

In 2009, the Global Nuclear Energy Partnership was cancelled by the Obama administration before it had progressed very far. Under Obama, the United States maintained a policy similar to that of the Clinton administration and pressured some partner nations, like the United Arab Emirates, to refrain from building their own enrichment and reprocessing facilities. The United States dropped plans to build reprocessing facilities at home but continued reprocessing research and development at a modest level.

The overall position of the Trump administration on reprocessing was unclear. The administration repeatedly proposed significant reductions in spending for nuclear fuel cycle research and development in areas including reprocessing technologies, but Congress restored much of that funding. However, the administration's director of the DOE Office of Nuclear Energy, Rita Baranwal, often spoke of her interest in spent fuel reprocessing and recycling, which she believes "can be better utilized to reduce the amount of nuclear waste over time" (Baranwal 2019). As of this writing, the Biden administration's position on these matters is unclear.

There are two primary considerations when assessing the terrorism and proliferation risks posed by a reprocessing technology—the attractiveness of the materials in the system and the difficulty of applying adequate safeguards and security.

MATERIAL ATTRACTIVENESS

Material attractiveness is related to the physical properties of the separated material. Could it be used to make a weapon directly, or would it need further refinement? And if it needs to be refined, how difficult and hazardous would it be to steal, transport, and convert it to a weapon-usable form? In other words, how attractive would the material be to someone seeking to build a bomb?

Conventional aqueous reprocessing—the PUREX process—separates plutonium from all the other elements in the spent fuel. Because separated plutonium can be used directly to make a nuclear weapon and is not highly radioactive (that is, not self-protecting), in sufficient quantity (2 kg or greater) it falls under Category I, the NRC’s highest category for physical protection.

Other proposed reprocessing technologies, such as pyroprocessing or alternative aqueous processes, would not produce a separate plutonium stream (if the process were operated as designed). Depending on the process, the plutonium would be mixed with combinations of other actinides, such as uranium, neptunium, americium, and curium, as well as certain fission products (primarily radioactive isotopes in the lanthanide series of the periodic table, some of which have similar chemical properties to actinides). The product would be somewhat more radioactive than separated plutonium due to the presence of some transuranic isotopes and fission products, arguably creating a deterrent to diversion and theft. Depending on the composition, some NLWRs might be able to incorporate this material directly without further processing, and it would retain its deterrent properties throughout the fuel cycle.

For years, the DOE argued that alternative reprocessing technologies such as pyroprocessing would have lower proliferation and terrorism risks than PUREX because their product would be more difficult to steal and process than separated plutonium (DOE 2006). However, by the end of 2008, the DOE had reassessed these technologies and concluded they would not be less risky than conventional reprocessing with respect to both nuclear proliferation and nuclear terrorism. In the December 2008 Draft Nonproliferation Impact Assessment for the Global Nuclear Energy Partnership,¹⁰ the DOE stated that “in the context of a potential diversion by a state, the nonproliferation benefits of blending plutonium either with uranium or the minor actinides are both very

modest” and that “the hazards of this level of radiation exposure by themselves would not prevent theft or malicious use of the material and would do little to deter someone who was willing to accept these risks” (DOE 2008).

The DOE’s draft nonproliferation impact assessment did find that retaining some lanthanide fission products in the reprocessing product would make it somewhat harder to handle safely. However, because the radiation levels would remain well below the “self-protection” threshold, it judged that the increased difficulty of safe handling would be marginal “and would be unlikely to deter an adversary who was willing to accept injury (or self-sacrifice)” (DOE 2008). That is, terrorists would be able to steal and chemically process the mixture to remove the lanthanides without being exposed to immediately life-threatening levels of radiation. As a result, the DOE concluded that all of the alternatives to PUREX it analyzed “involve materials that are sufficiently attractive for potential misuse that they require Category I physical protection measures” (DOE 2008). This assessment, based on studies conducted by nuclear weapons experts in the national laboratories, was consistent with the findings of the Union of Concerned Scientists report *Nuclear Power in a Warming World*, issued a year earlier (Gronlund, Lochbaum, and Lyman 2007) and other independent analysts (Kang and von Hippel 2005).

But even if the lanthanide fission products did make the end product harder to handle, there are strict limits on lanthanide impurities in fresh fast reactor fuel well below the level anticipated for the pyroprocessing product (Piet et al. 2010). The lanthanides would still have to be separated out, most likely using an aqueous process, before the product could be turned into fresh fuel—thus undoing any proliferation and terrorism-resistance benefits.

In sum, fuel cycles based on pyroprocessing or other advanced separation processes do not significantly reduce the material attractiveness of the reprocessing product relative to PUREX.

SAFEGUARDABILITY OF REPROCESSING AND FUEL FABRICATION PLANTS

The safeguardability of nuclear material processing facilities—or the ease of meeting safeguards goals for timely detection of diversion—is another major consideration in assessing the proliferation risks of closed fuel cycles. As discussed above, it is extremely challenging to monitor nuclear material effectively at bulk-handling facilities such as reprocessing and fuel fabrication plants, which process materials in forms such as liquids and powders. There are fundamental physical limits on the ability to keep track of weapon-usable material at industrial-scale facilities. These limits may make it extremely

difficult, if not impossible, for authorities to detect the diversion or covert theft of bomb-usable quantities of fissile material with enough warning time to prevent a country or terrorist from building a weapon.

In recent years, some have argued that this problem could be resolved through the practice of “safeguards by design”—the incorporation of design features into new facilities to improve the accounting of weapon-usable materials. However, while better facility design can help address the problem, it is unlikely to sufficiently mitigate the fundamental accounting problems (and associated diversion risks) at reprocessing plants, which are largely driven by the imprecision of available measurement techniques. This has become apparent in the recent DOE program known as MPACT (Material Protection, Accounting, and Control Technologies), which was intended, among other things, to demonstrate advanced safeguards by design principles for a model pyroprocessing plant, but was still unable to achieve acceptable detection probabilities for many diversion scenarios (Cipiti, Shoman, and Honnold, forthcoming). This is discussed further below.

MATERIAL ACCOUNTANCY IN PUREX REPROCESSING PLANTS

At a conventional PUREX reprocessing plant, spent nuclear fuel is processed in batches. It is first dissolved in acid, and the spent fuel solution is then piped to an input accountability tank, from which samples of the solution are drawn. The samples are taken to a laboratory for “destructive analysis,” where technicians chemically separate and purify the plutonium and uranium. This enables them to make the most accurate measurements of the total quantities of plutonium and uranium in the batch. (Even so, there are uncertainties resulting from sampling errors, measurement errors, and errors due to a lack of precise knowledge of the ratios of different plutonium isotopes.) Since methods for direct measurement of plutonium in spent fuel before it is dissolved are currently not precise enough to be useful in material accountancy, the first accurate measurement of the plutonium input occurs only after the fuel is dissolved and samples are taken.

At PUREX reprocessing and plutonium fuel fabrication plants, operators shut the facility down on a regular basis and remove as much material as possible from the process areas to measure the total amount of plutonium present. This physical inventory is necessary because it is not possible to accurately measure some of the material in processing areas. Even after the plant is flushed out, some plutonium—“residual holdup”—remains lodged in the equipment, so that the amount of plutonium coming out is generally less than the amount going in.

It is difficult to accurately account for the residual holdup. New plants are designed to have equipment in pro-

cess areas to measure holdup that cannot be feasibly removed and measured, but those *in situ* measurements have large uncertainties. This is why it is important to reduce the amount of residual holdup to as low a level as possible. However, over time, the residual holdup stuck in the plant can equal many bombs’ worth of plutonium. This makes it nearly impossible for operators and inspectors to determine with confidence that no plutonium has been covertly removed from the plant. Indeed, the facility MUF examples (“material unaccounted for”) described above demonstrate that the IAEA has not been able to meet its material accountancy goals at some PUREX and plutonium fuel fabrication facilities.

MATERIAL ACCOUNTANCY IN PYROPROCESSING PLANTS

Pyroprocessing plants have certain characteristics that would make them even harder to safeguard than PUREX reprocessing plants. In particular, it will be even more challenging at pyroprocessing plants to make the accurate measurements needed to keep track of plutonium and other weapon-usable materials (Mickum, McElroy, and Hertel 2014; Cipiti and Shoman 2018).

First, operators and safeguards inspectors cannot directly measure the quantity of plutonium going into a pyroprocessing plant. In these plants, metallic spent fuel is placed in a basket and immersed in an electrorefiner vessel filled with molten salt. As the spent fuel dissolves into the molten salt, plutonium is distributed to different parts of the system in an inhomogeneous manner. There is no counterpart to the input accountability tank in PUREX plants, where all the plutonium in a batch is first contained and representative samples can be taken and measured before further processing. Without an accurate measurement of how much plutonium is going into the facility, it is difficult to know when a significant quantity of material has gone missing. This would not be a problem if it were possible to accurately calculate the quantity of plutonium and other actinides in the initial spent fuel. However, such calculations typically have uncertainties on the order of 10 percent, far too large an error to be useful in material accountancy.

Second, pyroprocessing is an inherently impure separation process, which further hinders the ability to directly and accurately measure all of the plutonium in the system. The process is designed to separate most of the uranium from the remainder of the spent fuel. But the separated uranium plates out on a steel cathode in the form of hard deposits called dendrites. In order to collect the deposits, a hammer is used to chip away at the dendrites, leaving a considerable amount of uranium stuck to the cathode—anywhere from two to 10 percent. Because the uranium product is also contaminated with a significant amount of plutonium, this residue also contributes to the plutonium measurement uncertainty.

Third, an essential element of reprocessing plant safeguards—periodically cleaning out the plant and taking a physical inventory of the material—cannot be performed in a timely manner at pyroprocessing plants. At these plants it is necessary to build up and maintain a minimum concentration of plutonium and other actinides in the molten salt in order for the process to work (Cipiti et al. 2012). There is no need from an operational perspective to process and purify the salt until the actinides have accumulated to the extent that they become a safety concern—for instance, if the risk of an accidental chain reaction (a criticality event) becomes too high. (At the pyroprocessing facility at the Idaho National Laboratory, for example, because most of the spent fuel that has been pyroprocessed contained very little plutonium and other TRU, the salt has remained in the vessel without being cleaned up since the process was started up in 1996.) Thus, cleaning out the salt to directly assay its nuclear material content frequently enough to meet safeguards goals could disrupt operation of the plant. PUREX plants have no counterpart to the electrorefiner salt that would cause an analogous problem.

NON-DESTRUCTIVE ASSAY

Because operators would not clean out pyroprocessing plants frequently, it would not be possible to directly measure the plutonium accumulating in the plant using destructive analysis. One way to address this problem would be to directly measure the plutonium in the plant by non-destructive assay methods, such as counting the neutrons that fissile isotopes emit through processes such as spontaneous fission.

However, this approach could at best indirectly—and imprecisely—determine the amount of plutonium. Neutron counters are designed to detect neutrons emitted by radioactive isotopes such as plutonium-239, but it is hard to distinguish the neutrons emitted by one isotope from those emitted by another. Therefore, if plutonium is mixed with other neutron-emitting materials, neutron counters are not very good at identifying the specific isotopes. This is particularly problematic in pyroprocessing plants because the separation process is designed to keep plutonium mixed together with other neutron-emitting TRU in spent fuel, including curium isotopes. However, the curium isotopes Cm-244 and Cm-242 emit neutrons at high rates that swamp the neutron emissions from the plutonium isotopes; neutron counting thus can only directly measure the amount of curium.

The amount of plutonium can be determined indirectly if the ratio of plutonium to curium is known throughout the system. But this will be the case only if this ratio can be accurately determined in the original spent fuel and if it remains

constant throughout the entire process—two questionable assumptions.

Some researchers developing pyroprocessing safeguards rely heavily on presumed knowledge of the plutonium-curium ratio in spent fuel. However, there are numerous technical problems with the approach. For example, the ratio can only be estimated by computer simulations of reactor operations that have unacceptably high uncertainties. According to Oak Ridge National Laboratory researchers, “the Pu/Cm ratio may not allow definitive safeguards conclusions to be drawn because the Pu material unaccounted for (MUF) may exceed the significant quantity of 8 kg” (Mickum, McElroy, and Hertel 2014).

And modeling by Sandia National Laboratory researchers has shown that for a small plant (100 metric tons per year) uncertainties of 0.5 percent for the TRU mass in the pyroprocessing cell and 1.0 percent for input and output measurements would be necessary to meet IAEA goals, including detection of a protracted diversion of 8 kg of plutonium within one year. These researchers point out that those uncertainty targets “may be difficult to achieve” by using non-destructive assay (Cipiti et al. 2012).

More recent results of the DOE MPACT (Material Protection, Accounting, and Control Technologies) study mentioned above are consistent with these findings (Cipiti, Shoman, and Honnold, forthcoming). The study analyzed the effectiveness of safeguards measures for detecting various diversion scenarios at a pyroprocessing plant with a capacity of 100 metric tons per year that incorporates safeguards-by-design principles. The study found that to achieve at least a 95 percent probability of detection of a diversion of 8 kg of plutonium in 30 days for all scenarios, all key measurements would have to have an uncertainty of 1 percent, and the facility would have to be shut down every eight days to conduct a physical inventory. For more realistic uncertainties of 5 percent, detection probabilities for different diversion scenarios would range from 63 percent to as low as 13 percent, falling far short of the 95 percent goal (Cipiti, Shoman, and Honnold, forthcoming). The study found even worse results if the objective were to detect the diversion of 2 kg of Pu in seven days, the NRC’s regulatory goal. For larger, commercial-scale pyroprocessing plants, the difficulty of detecting diversions would be even greater.

A recent survey of a wide range of potential non-destructive assay measurement techniques at pyroprocessing plants found few possibilities with theoretical uncertainties of less than 1 percent; most had uncertainties between 1 and 15 percent. The more precise techniques would require pure samples and/or take hours or weeks to obtain measurements, and would not be useful by themselves for timely detection

of diversions (Coble et al. 2020). The study concluded that individual state-of-the-art measurement approaches are insufficiently precise to meet the necessary requirements for achieving material accountancy goals.

CONTAINMENT/SURVEILLANCE AND PROCESS MONITORING IN PYROPROCESSING SAFEGUARDS

Given the large measurement uncertainties to be expected at pyroprocessing plants, the IAEA would have to rely heavily on complementary measures to meet its inspection goals, including containment and surveillance. However, these are inadequate substitutes for accurate material accounting. For instance, if surveillance such as closed-circuit television coverage were interrupted, then inspectors would not be able to rule out the possibility that material was diverted during the outage. The only way to resolve this problem would be to conduct an inventory to verify that no material is missing. But if the material accounting system has large uncertainties, then it may take a very long time, or may not even be possible, to verify that a significant quantity of material was not diverted during the loss of the surveillance system.

An additional complementary method being studied by researchers utilizes process monitoring. This is a qualitative approach to identify deviations from normal process parameters that would alert operators if a diversion of material or another type of abnormal event were taking place (Cipiti and Shoman 2018). However, such approaches themselves have major limitations, including the potential for a high rate of false alarms. As with containment and surveillance measures, their use would not obviate the need for precise quantitative techniques to quickly determine whether a diversion had occurred should a process anomaly be detected.

To summarize, pyroprocessing plants would likely present greater proliferation risks than conventional PUREX plants because they would be more difficult to safeguard. The IAEA is already unable to meet its material accountancy goals at PUREX plants using the best available technologies, and these would be less effective or not usable at all at pyroprocessing plants. Moreover, there are no techniques currently available or on the near horizon that could accurately

measure the amount of plutonium going in, the amount of plutonium going out, or the amount of plutonium within a pyroprocessing plant. As a result, nuclear plant operators and safeguards inspectors would have even more difficulty detecting diversions of weapon-usable material in a timely manner than at PUREX plants.

The lack of an effective safeguards approach for NLWRs and their fuel cycles may be an obstacle to their deployment in non-nuclear weapon states such as Canada, because the IAEA will need to approve a safeguards approach before such reactors could operate. For example, the Canadian Nuclear Safety Commission is now undertaking pre-licensing activities for a number of NLWR designs that may involve on-site reprocessing or have other features that would require new safeguards techniques. The lack of a safeguards approach will be one of a number of significant obstacles to the rapid deployment of such reactors in Canada that some vendors hope to achieve.

Nuclear proliferation is not a risk in the United States simply because it already possesses nuclear weapons and is designated as a nuclear-weapon state under the Nuclear Non-Proliferation Treaty. As such, it is not obligated to submit its nuclear facilities and materials for verification by the IAEA, although it is free to do so on a voluntary basis. However, nonproliferation is relevant to US reactor development both because US vendors seek to export new reactors to other countries and because other countries are likely to emulate the US program. The United States has the responsibility to set a good international example by ensuring its own nuclear enterprise meets the highest nonproliferation standards.

One way to do that would be for the United States to designate all new nuclear reactors and fuel-cycle facilities as eligible for IAEA safeguards under its voluntary offer agreement with the IAEA. This would give the IAEA an opportunity to develop verification approaches for new types of facilities—if such approaches are feasible. Unfortunately, there is no indication that the United States is planning to make any of its proposed new NLWR projects eligible for IAEA safeguards.

Liquid Sodium–Cooled Fast Reactors

In 2017 the Department of Energy (DOE) identified liquid sodium–cooled fast reactors as one of two non-light-water reactor (NLWR) technologies that it believed were sufficiently mature to support construction of either a test reactor or a commercial demonstration reactor in the “near future” (the other being the high-temperature gas–cooled reactor discussed in chapter 6) (Petti et al 2017). Today, the DOE is moving forward with plans to build both a fast test reactor *and* a commercial demonstration fast power reactor—both based on the General Electric–Hitachi (GEH) PRISM sodium-cooled, metal alloy–fueled design.

This chapter discusses the history and current status of sodium-cooled fast reactors, safety issues, and the time scale, costs, and risks of building a commercial-scale demonstration PRISM reactor and its associated fuel cycle facilities, which may include a facility to pyroprocess its spent fuel. It assesses the steps needed before a reactor vendor would be ready to build a commercial-scale demonstration fast reactor or test reactor based on the PRISM design.

In September 2020, the DOE decided to proceed with engineering design of the 300 megawatt-thermal (MWth) Versatile Test Reactor (VTR), which would most likely be built at the Idaho National Laboratory (INL). Also, in October 2020 it selected the 840 MWth (345 megawatt-electric (MWe)) TerraPower–GEH Natrium reactor as one of two designs to be built under its Advanced Reactor Demonstration Program (ARDP). If DOE decides to proceed with construction of the VTR, it anticipates the reactor will be operational between 2026 and 2031—as soon as eight years after the project began in 2018 and four years after construction begins in 2022 (INL n.d.). The Natrium is supposed to be operational by

2025–2027, according to the terms of the ARDP. A third metal-fueled fast reactor project with some similarities to the PRISM design is also underway: the Oklo, Inc. Aurora 1.5 MWe micro-reactor. Oklo is applying for a combined operating license to build an Aurora unit at the INL, which it anticipates could be operational by the early- to mid-2020s.

The DOE judges that PRISM is a mature reactor design, and the VTR and Natrium projects are proceeding on the expectation that both can skip performance demonstrations because prior fast reactor demonstrations have provided the necessary data.¹¹

However, as discussed below, the PRISM design has never had a full-scale performance demonstration: the VTR and the Natrium will serve as the first large-scale demonstrations of PRISM technology. It is far from clear that prior fast reactor experience has provided adequate supporting evidence that full-scale PRISM reactors can be operated safely or reliably. Thus proceeding with construction of the VTR and the Natrium without conducting prototype testing could pose unacceptable risks to public health, safety, and security, as well as to the success of either project.

History and Current Status

The fast reactor—a nuclear reactor that does not require a moderator material to slow down fission neutrons—is an old technology. The concept of the fast breeder reactor was originally conceived by Leo Szilard and other Manhattan Project scientists in 1944. The first nuclear reactor to generate electricity in the world was the Experimental Breeder Reactor-I (EBR-I), which famously lit four light bulbs in December

1951. The EBR-I also has the distinction of being the first US nuclear reactor to experience an unplanned core melt, in November 1955.¹²

Since water cannot be present in a fast reactor core because it would slow down the neutrons, it is necessary to use a heavier substance as a coolant, such as a liquid metal. The US EBR-I used a sodium-potassium coolant, and the Soviet BR-2 used a mercury coolant. However, based on extensive testing of coolants by the United States and other countries, all other fast reactors built over the last 50 years have used sodium. Other potential candidate liquid metal coolants are molten lead or lead-bismuth alloy, but lead-bismuth so far was used only in a small number of Soviet nuclear submarines, resulting in three deadly accidents.

Since the 1950s, there has been considerable research and development of sodium-cooled fast reactor technology around the world (see Table 1, p. 17). There are five such reactors operating today: four experimental demonstration reactors in India, Russia, and China, and one commercial demonstration reactor in Russia. India's 500 MWe demonstration reactor has been delayed for more than a decade and is currently slated to begin operation in December 2021 (WNA 2021). While there are new types of fast reactors under development in several countries, there are no other sodium-cooled fast reactors currently being built.

The only large fast reactors that have been connected to the electricity grid that are sodium-cooled are the French Superphénix reactor and the Russian BN-600 and BN-800 reactors. Superphénix, which operated for 13 years, was shut down more than half of the time for repairs (Cochran et al. 2010).

RECENT DEVELOPMENTS

Two sodium-cooled fast reactor concepts were submitted to the DOE for evaluation as potential demonstration reactors in its 2017 study: the GEH PRISM modular reactor concept and Argonne National Laboratory's AFR-100. The DOE concluded that a commercial-scale demonstration PRISM reactor was ready to be built because the design was based on the Experimental Breeder Reactor-II (EBR-II), a small test reactor in Idaho that operated for about three decades. Each PRISM module would produce 471 MWth or 165 MWe, about eight times greater than the EBR-II. (A larger model would produce 311 MWe.) In contrast, the DOE concluded that the AFR-100, which would produce 100 MWe, was sufficiently different from the EBR-II that it would require validation at an experimental scale before a commercial-scale demonstration plant could be built.

A number of startup companies are pursuing commercialization of sodium-cooled fast reactors in the United States and abroad. These include TerraPower, the company financed

in part by Bill Gates, the ARC-100, and Oklo, Inc (actually potassium-cooled, with sodium-bonded fuel). All three companies are developing reactors that are based to varying extents on the GE-Hitachi PRISM reactor design, which itself is an evolution of the EBR-II.

TerraPower, likely the best-capitalized of the startups, was founded to develop a once-through traveling-wave "breed-and-burn" reactor (see chapter 8), but it is currently focusing on the more conventional Natrium fast reactor. Initially, the company has planned to build a 600 MWe reduced-scale prototype traveling-wave reactor in China by as soon as 2022, although it would not have been capable of breed-and-burn operation because some of the necessary technologies—including ultra-high burnup fuels—have not yet been developed. TerraPower has now rebranded this more conventional design as the Natrium, which the DOE has selected for deployment by 2027 under the ARDP. At 100 MWe, the ARC-100, which is being developed by Advanced Reactor Concepts (ARC), a private company founded by former EBR-II engineers in collaboration with GE-Hitachi. In July 2018 ARC signed an agreement with New Brunswick Power in Canada to explore deployment of the reactor at the Point Lepreau nuclear plant. And Oklo, a 1.5 MWe "micro-reactor," submitted a combined operating license application to the Nuclear Regulatory Commission (NRC) in February 2020 for construction of a single unit at the INL. Oklo has said it would like to deploy its first reactor by the "very early 2020s" (DeWitte 2016).

Also, as discussed above, with \$65 million in support from Congress in fiscal year (FY) 2020 and \$45 million in FY 2021, the DOE is proceeding with detailed design of the VTR, which it hopes to build and operate as early as 2026, pending a final go-ahead in 2022. The ostensible purpose of the VTR is to produce a high flux of fast neutrons for assisting in the development of fuels and materials for fast reactors. But in order to do that, the VTR itself will be a moderately sized, 300 MWth sodium-cooled fast reactor based on the GEH PRISM design. Despite being labeled as a "test" reactor, the VTR would be larger than the largest materials-test reactor in the world, the 250 MWth Advanced Test Reactor at the INL, and about as large as the PRISM reactor design that GEH had previously submitted to the DOE as a candidate demonstration reactor. In addition, as discussed below, the VTR would have novel characteristics that have not been sufficiently demonstrated in previous fast reactors (Lyman 2018b). Therefore, it would be more accurate to characterize the VTR as a demonstration reactor rather than a materials test reactor. However, unlike the Natrium, since the VTR will not be generating electricity, it will not demonstrate this key aspect of commercial power operation.

Other countries are deferring long-planned sodium-cooled fast reactor projects. France, a long-time proponent

of fast reactor technology, decided in 2019 to postpone plans to build a demonstration fast-breeder reactor called ASTRID until the “second half of the century,” effectively terminating the project after spending over \$800 million on it (Patel 2019). Also, Russia deferred construction of its BN-1200 fast reactor—a long-planned next step in commercialization of the technology—until after 2035. Russia is continuing to pursue construction of a nuclear reactor complex, called the “Breakthrough” project, that will include a 300 MWe lead-cooled fast reactor called the BREST-300. However, the DOE believes that fast reactors using lead-based coolants are less mature than sodium-cooled fast reactors and would require a larger effort to commercialize them.

FAST REACTOR FUEL TYPES: METALS AND OXIDES

One important distinction among fast reactors is the type of fuel that they use. The two most common fast reactor fuels are metals and oxides, although other compounds such as nitrides have also been pursued. The choice of fuel type seems to be based more on institutional and national preference than on any definitive technical factors. For its fast reactors, the United States has used both oxides and metals in the past but currently prefers metal, while France has chosen oxide, and Russia is developing nitrides.

The early fast reactors (EBR-I, EBR-II, Fermi-1, and Dounreay Fast Reactor) used a metal fuel, in part to maximize the potential for breeding plutonium (see chapter 3). Compared to compounds such as oxides, metal fuels are denser, allowing for higher concentrations of neutrons. Also, the energy of the neutrons is higher (since there are no lighter elements, such as oxygen, that can slow down neutrons). Both of those factors improve the breeding potential of the reactor. However, radiation causes the metal fuel to swell over time and potentially break through its cladding if used in the reactor for more than a short period of time. Since commercial reactors have an economic incentive to use the fuel for a relatively long time, this issue (as well as other safety concerns) led to the development of ceramic oxide fuel, which subsequently was adopted for most fast reactor projects around the world.

However, metal fast reactor fuel still has its advocates in the United States and a few other nations including South Korea. As noted above, a number of fast reactors under development in the United States, including the VTR, the Natrium, the ARC-100, and Oklo’s Aurora, are designed to have metal fuel. Researchers have partly addressed the clad failure problem by providing extra space to accommodate fuel swelling and other modifications. However, significant further development is needed to resolve other issues that limit fuel burnup, such as the high pressure from fission product gases released from the fuel.

Another gap in the experimental record for metal fuels is the lack of data on plutonium-based fuels. Most fast reactors under development in the United States would likely need to use fuels fabricated with plutonium for breeders or plutonium and other transuranic elements (TRU) for burners. The DOE’s choice of fuel for the VTR is an alloy of plutonium and low-enriched uranium metal. However, most plutonium fuels irradiated in fast reactors around the world have been oxides. In the United States, other than a small fraction of test fuel elements (less than 0.5 percent), the metal fuels irradiated at fast reactors contained uranium at the outset (either highly enriched or depleted) and not plutonium. This is an issue because plutonium and other TRU have different physical, chemical, and nuclear properties than uranium, which will lead to differences in the performance of these fuels in reactors.

Fast Reactors: Cost Considerations

This report does not attempt to conduct a comprehensive analysis of the costs of NLWRs. Developing cost estimates for any new nuclear power technology is a treacherous business, even for LWR designs that are close cousins of the operating fleet. The total cost of the AP1000 LWR project at the Vogtle nuclear plant in Georgia is now projected to be \$28 billion, twice the original estimated cost (Nuclear Engineering International 2020).

However, one relatively safe bet is that sodium-cooled fast reactors will be significantly more expensive to build and operate than LWRs. There are fundamental technical reasons why this is the case. This has also been borne out by the historical experience with sodium-cooled fast reactors, for which there are many more examples than with other types of NLWRs.

ADDITIONAL SYSTEMS, SPECIALIZED EQUIPMENT, AND MORE ROBUST CONSTRUCTION

Liquid metal-cooled fast reactors have higher capital costs than LWRs because they require many systems that LWRs do not need (Zhang et al. 2009). For instance, liquid-sodium cooled reactors typically have an additional, intermediate coolant loop that transfers heat from the primary sodium coolant system to the steam generators. This system acts as a buffer between the radioactive sodium in the primary coolant system and the water in the steam generators, avoiding the potential for a sodium-water explosion that could disperse radioactivity. The reactor also must have prevention, detection, and mitigation systems for sodium leaks (Zhang et al. 2009). Fast reactor vessels would also cost more than LWR vessels. Although their walls do not need to be as thick

because the system pressure is lower, they are also larger and more complex, containing additional structures and requiring specialized equipment (Braun 2012). Lastly, fast reactor containment structures would need to be larger and more robust than LWR containments if regulators ultimately require them to withstand the very high pressures and temperatures of severe accidents, including the so-called hypothetical core disruptive accident discussed below.

Past studies of fast reactors have estimated that their capital costs would exceed those of LWRs by 25 to 75 percent (Cochran et al. 2010). One comprehensive survey assumes a fast reactor premium of 20 percent for both overnight capital cost (e.g., without considering interest during construction) and operating and maintenance costs, although it does not provide the basis for this assumption (NEA 2013). A review of the historical experience of building demonstration-scale sodium-cooled fast reactors has found that their capital costs have typically been more than twice those of contemporary LWRs (Cochran et al. 2010).

Given the problems caused by liquid sodium, could fast reactors be cheaper if they used a different coolant? Some argue that because other coolants such as lead or lead-bismuth do not react violently with water, fast reactors using them would not need the costly intermediate heat exchanger that sodium-cooled reactors require. However, the neutron activation of bismuth generates polonium-210, a hazardous radioisotope, so the intermediate heat exchanger may still be necessary in lead-bismuth-cooled plants to isolate the radioactive coolant from the steam generators and protect workers from excessive radiation exposure. And lead's chemical toxicity and extreme corrosivity would complicate the management of large volumes of the molten metal.

DISECONOMIES OF SCALE IN MODULAR CONSTRUCTION

Another issue affecting the capital costs of fast reactors is a potential safety limit on the power rating of a single unit. While LWRs in principle can achieve greater economies of scale by getting larger without necessarily compromising safety, fast reactors become less safe as they get larger because the sodium void coefficient tends to increase with reactor size. Therefore, a safer way to build large baseload fast reactor plants would be to construct multiple modules rather than to build a single large reactor. But this would tend to increase cost due to diseconomies of scale. Although proponents of small modular reactors claim that this cost penalty would be outweighed by the efficiencies gained from mass production of multiple modules, there is no compelling evidence at present to support this assertion (Lyman 2013). Also, smaller fast reactors tend to leak more neutrons and have worse sustainability performance (see chapter 3).

HIGHER OPERATING AND MAINTENANCE COSTS

Operating and maintenance costs for fast reactors would also be greater than for LWRs. Liquid sodium is a difficult material to work with in a number of ways. For example, because it is opaque, safety inspections are more difficult for reactor structures immersed in sodium.

Additional security and material accountancy measures will also increase the operating cost relative to LWRs. Because of sodium's opacity and chemical reactivity, it is difficult for safeguards inspectors to keep track of fast reactor fuel within the reactor and intermediate storage pools containing spent fuel. Fast reactors using plutonium fuels are generally Category I facilities, which are required in the United States to maintain armed response forces capable of preventing both sabotage and theft of weapon-usable materials. In contrast, armed security forces at LWRs need only to protect against sabotage, since their LEU fuel is not directly usable for weapons. Security measures for protection against theft of weapon-usable materials are generally more stringent—and costly—than those for protection against sabotage.

Some fast reactor advocates are keenly aware that the additional costs associated with this technology would make nuclear power less economical than for the current fleet of LWRs, which is already struggling to compete with low-cost natural gas-fired generation and wind and solar power. But some claim that they can build fast reactors that not only will be cheaper than LWRs but will be competitive with natural gas. In Russia, Rosatom has initiated a project called “PRORYV” (Breakthrough), with a primary goal to establish the competitiveness of the nuclear power industry, which project leaders believe has been in crisis for the last 30 years (Adamov et al. 2016). PRORYV is undertaking an effort to develop fast reactors with capital costs 20 percent below those of LWRs. However, for the reasons discussed above, doing so may require significant compromises in safety and security, and there is good reason to be skeptical.

REPROCESSING AND RECYCLED FUEL COST

In addition to the greater capital and operating costs of fast reactors, there is also the cost of the fuel cycle. As discussed in chapter 3, most fast reactors can only realize their full potential for increasing sustainability in a closed fuel cycle with recycling of plutonium in fresh fuel.

Plutonium-based fuel will be considerably more expensive than LEU fuel for LWRs. For fast reactor designs using plutonium and possibly other TRU elements, such as neptunium and americium, fuel production would involve some type of spent fuel reprocessing. Reprocessing is an extremely costly industrial enterprise that requires the construction

and operation of high-capital-cost, heavily shielded facilities with remote-controlled equipment, as well as many safety and security features necessary to manage separated plutonium and highly radioactive fission products. Reprocessing plants convert a single waste form—spent nuclear fuel—into multiple waste streams, which pose additional challenges for management and disposal. In addition, fabrication plants for fresh reactor fuel containing plutonium or other TRU elements require additional safety, security, and waste management measures compared to LEU plants, adding to the cost.

Many studies over the last few decades have confirmed that nuclear fuel cycles including reprocessing and plutonium fuel fabrication will increase cost relative to the once-through cycle with direct disposal of LEU spent fuel that is used by LWRs (Bunn et al. 2003; MIT 2011; NEA 2013). A review by the Nuclear Energy Agency reported that such studies found cost premiums for the fast reactor–based closed fuel cycle ranging from 25 to 42 percent (NEA 2013). However, reprocessing advocates argue that this would lead to only a small increase in the total cost of electricity, since the fuel-related cost is only a fraction (less than 20 percent) of the electricity cost. Nevertheless, the absolute cost premium that would be borne by taxpayers or electricity consumers could be on the order of many tens of billions of dollars over the lifetimes of all of the facilities needed to research, develop, and implement the respective fuel cycles.

These estimates find that closed fuel cycles are more expensive even after accounting for several factors that tend to offset the additional costs of reprocessing and recycle. These include the reduced demand for natural uranium and the potential to reduce the required footprint for geologic repositories for long-lived radioactive wastes.

In any event, the potential impact on the total fuel cycle cost of these and other factors is highly sensitive to input parameters that have large uncertainties. For instance, analyses have shown that the most important parameter in determining the relative costs of different fuel cycles is the price of uranium. The cost savings resulting from a reduced need for natural uranium is only significant if uranium is expensive. A Nuclear Energy Agency study found that a fuel cycle in which spent fuel was repeatedly reprocessed and the plutonium used in both LWRs and fast reactors would only become economically attractive for uranium prices of \$270 to \$300 per kilogram—nearly 4 times the February 2021 spot price of around \$77 per kilogram (NEA 2013).¹³ The second most important factor was the assumed cost of reprocessing. The results were also very sensitive to the assumed cost premium for fast reactors relative to LWRs. On the other hand, the calculated fuel cycle costs were far less sensitive

to the assumed cost of a geologic repository—which means that even if adoption of a closed fuel cycle did reduce the need for geologic repository capacity, it would not translate into significant cost savings. This calls into question the real value of one of the major selling points of closed fuel cycles (see chapter 3).

The Nuclear Energy Agency study asserts that the estimated difference in the cost of closed and open fuel cycles is small enough that it is washed out by the uncertainties in the input parameters. However, for some choices of those parameters the magnitude of the cost difference itself could be far larger. For example, the study chose a uranium price of \$130 per kilogram as its base case value (corresponding to the spot price in early 2011) and expected the price to rise in the future. But since then the price has plummeted to 60 percent of that amount, increasing the cost premium of plutonium fuel. Also, the total cost of reprocessing used in the study—a critical parameter—ranged from \$579 to \$2,640 per kilogram, depending on numerous assumptions; however, this may underestimate the actual cost of fast reactor fuel reprocessing systems.

Proponents of pyroprocessing, a key component of several proposed advanced-reactor fuel cycles, argue that the technology would be cheaper than conventional aqueous reprocessing. They have a long way to go to demonstrate that, however. To date, the actual cost of the only operating pyroprocessing system has averaged more than \$50,000 per kilogram of spent fuel—20 times greater than the highest value assumed in the Nuclear Energy Agency study.

In summary, recent studies have confirmed that the adoption of a closed fuel cycle utilizing fast reactors and reprocessing will increase the cost of nuclear power. Given that reprocessing and plutonium recycling make waste management more difficult while simultaneously increasing cost and safety and security risks, it is hard to see the benefits of advanced reactor systems that are lauded for their ability to “consume” nuclear waste.

Safety

Sodium-cooled fast reactors have inherent safety disadvantages relative to LWRs. Fast reactor designers have worked for decades to address these issues, but for the most part they have failed to resolve them. One of the primary concerns is that compared to LWRs, these reactors commonly have a fundamental and significant instability: a positive void coefficient. As discussed in chapter 2, this means that if the temperature of the sodium coolant increases and the sodium boils, the power of the reactor typically increases. This positive feedback effect could lead to a rapid increase in pressure and

temperature, further coolant boiling, and core damage. This effect was a major contributing factor to the 1986 Chernobyl accident, which involved a type of reactor that had a positive void coefficient. It has proven very difficult to design fast power reactors that are entirely free of this problem. In contrast, LWRs typically have a negative void coefficient and exhibit more stable behavior: as the power increases and the coolant (water) heats up and becomes less dense, the reactor power decreases.

The properties of the different chemical forms of fast reactor fuels, such as metals, oxides, or nitrides, can affect reactors' safety and performance. For instance, advocates of metal-fueled fast reactor designs, such as PRISM, claim they are inherently safe because the metal fuel expands more rapidly than ceramic oxide fuel when heated, causing negative reactivity feedback that would reduce the reactor's power production even if the sodium boils. However, this claim that metal-fueled fast reactors are inherently safe, despite the presence of a positive sodium void coefficient, is overblown and misleading, as discussed below.

As discussed in chapter 1, one of the original objectives of the DOE's Generation IV program was to develop NLWRs that are significantly safer than current-generation LWRs. However, given the significant uncertainties and unresolved safety issues of sodium-cooled fast reactors, it remains far from clear that they can meet this objective. In a 2015 review, the French nuclear safety research organization Institut de Radioprotection et de Sûreté Nucléaire (IRSN) stated that it could not determine, "in view of design differences and the current state of knowledge and research," whether sodium-cooled fast reactors would be significantly safer than LWRs currently under construction such as the EPR (IRSN 2015).

Moreover, sodium-cooled fast reactors have a number of characteristics that may render them less safe than current-generation LWRs. Although reactor designers have been aware of most of these problems since the early days of the technology, the problems have proven difficult to resolve. In addition to the positive void reactivity problem, others include the use of chemically reactive liquid sodium coolant; the potential for rapid, hard-to-control power increases; and even the possibility of a small nuclear explosion, or as has often been referred to euphemistically, an "energetic core disassembly."

SODIUM COOLANT SAFETY ISSUES

Liquid sodium coolant has several characteristics that appear—initially—to provide safety advantages compared to water. Sodium has a high boiling point of nearly 900°C and does not need to be kept under high pressure during reactor opera-

tion. It also does not corrode reactor structures at normal operating temperatures.

On the other hand, sodium is a highly reactive material that combusts upon contact with air and reacts violently with water. Problems resulting from leaks of liquid sodium coolant have played a significant role in the poor performance of fast reactor demonstration projects around the world. For example, the Monju facility in Japan was shut down for more than two decades after experiencing a sodium fire in 1995 and is now to be decommissioned.

Even though liquid sodium exerts low pressure during normal operation in a fast reactor, the increases in pressure and temperature resulting from a sodium fire could be severe and potentially breach the reactor vessel, piping, and containment. To reduce the risk of the sodium fires that have affected many fast reactor projects, current designs are equipped with elaborate systems for sodium leak detection, leak mitigation, and fire suppression. In addition, unlike LWRs, sodium-cooled fast reactors must have an intermediate sodium coolant loop between the primary system and the steam production system, to reduce the risk of radioactive sodium in the primary coolant system coming into contact with—and violently reacting with—water. As mentioned above, to avoid the difficulties of liquid sodium, some have proposed using a lead-bismuth or pure lead coolant instead, which would eliminate the risk of violent sodium-air or sodium-water reactions. However, stainless steel is highly vulnerable to corrosion from molten lead, introducing other problems.

REACTIVITY FEEDBACKS IN FAST REACTORS

As discussed in chapter 2, one of the classes of events that could cause a severe reactor accident is a rapid increase in power resulting from a runaway chain reaction—the cause of the 1986 Chernobyl accident. Some types of nuclear reactors have a property that significantly reduces the likelihood or severity of this type of accident: an inherent tendency to slow down the fission process if the fission rate (and temperature) increases (see chapter 2). Indeed, for all US reactors, the NRC's General Design Criterion 11 requires that "the reactor core and coolant system be designed so that in the power operating range, the net effect of prompt inherent nuclear feedback characteristics tends to compensate for rapid increases in reactivity" (Appendix A to 10 C.F.R. §50 (1971)).¹⁴

The reactivity feedback behavior of fast reactors differs significantly from that of LWRs (see Box 6, p. 61). This is quantified by differences in their reactivity coefficients, which describe how the reactivity of the system changes in

continued on p. 63

Reactivity Effects in Fast Reactors

POWER COEFFICIENT

For an assessment of the stability of a reactor system, all of the various effects that change with reactor power or temperature must be considered together. If the overall reactivity decreases with an increase in the reactor power, the reactor is inherently stable.

While LWRs generally have negative moderator, fuel, and void temperature coefficients—and an overall negative reactivity feedback with respect to increases in temperature and power—the situation for metal-fueled fast reactors is more complicated. They may have positive coolant temperature and void coefficients, but a negative fuel coefficient. The magnitude of these coefficients, as well as those associated with other feedback effects, will determine the overall stability with respect to changes in power. These analyses are quite complex and rely on a combination of often-sparse experimental data and large-scale calculations. The uncertainties in such analyses may be large, making it hard to accurately predict the reactor's behavior.

One challenge in designing sodium-cooled fast reactors is that making changes to reduce the sodium void coefficient to reduce the severity of a sodium boiling accident can increase the severity of a “transient overpower accident,” a rapid increase in power that could be caused by the ejection of a control rod. A reactor core that readily leaks neutrons would require greater excess reactivity in the fuel, meaning that the control rods would have to be stronger neutron absorbers (have higher “worth”) to maintain the necessary power level at the beginning of the reactor cycle. But the presence of higher-worth control rods would increase the severity of a rod ejection event. Addressing these safety concerns simultaneously has proven difficult, especially for metal fuels (Van Tuyle et al. 1992).

FUEL TEMPERATURE COEFFICIENT

The fuel temperature coefficient is a measure of how the reactivity changes with the fuel temperature. As discussed in chapter 2, one important phenomenon that affects the fuel temperature coefficient is the Doppler effect. As the fuel temperature increases, the U-238 in the fuel absorbs more neutrons but does not fission, thus slowing down the fission process. The time scale over which the feedback occurs is nearly instantaneous, because it is a response to an increase in motion of the nuclei within the fuel. The magnitude of this effect depends on the properties of the fuel and the neutron speeds within the reactor.

For the low-enriched uranium oxide fuel used in LWRs, the fuel temperature coefficient of reactivity is negative: the reactivity decreases if the reactor temperature increases (which could be caused, for example, by an increase in power). This is because the fuel is nearly 90 percent U-238, and the neutron speeds are in the range where they are most susceptible to being absorbed.

Compared to in LWRs, the Doppler effect in fast reactors is less effective because there is less absorption in U-238 at high neutron speeds. And fast reactors that use metal fuels have even smaller Doppler feedback than those with oxide fuels. This is because with the absence of oxygen in the fuel, which has a slight moderating effect, neutrons will have higher average neutron speeds.

Another important phenomenon that affects the fuel temperature coefficient in fast reactors is thermal expansion of the fuel, which can provide a relatively rapid negative reactivity temperature feedback effect and help to stabilize the reactor power. Both oxide and metal fuels will expand as they get hotter, but the expansion is greater for metal fuel. This expansion would primarily take place along the direction of the fuel rod (the axial direction). As the fuel expands, it becomes less dense, reducing the chance that a neutron will strike a nucleus and cause it to fission, thereby reducing reactivity.

Proponents of metal-fueled fast reactors such as PRISM highlight the negative reactivity effect of fuel expansion as a major passive safety feature. For example, the PRISM website states that “in the event of a worst-case-scenario accident, the metallic core expands as the temperature rises, and its density decreases slowing the fission reaction. The reactor simply shuts itself down” (GEH 2021). However, this statement is misleading. First, the thermal expansion effect in metal fuels merely compensates for the much smaller Doppler effect relative to oxide fuels. The inherent prompt feedback of metal fast reactor fuels is no greater than that of LWR fuels. Second, the passive feedback would not by itself “shut [the reactor] down”—that is bring the reactor to a zero-power, subcritical state (NRC 1994). To make that happen, power plant operators would have to activate shutdown systems, such as control rods. (This is true for most reactors.) Finally, negative feedback associated with core expansion is not a feature of metal-fueled reactors only. Although metal fuels expand more than oxides as they heat up, analysis has shown that the actual negative temperature feedbacks from core expansion (both axial and radial) are comparable in metal-fueled and

CONTINUED ON P. 62

oxide-fueled fast reactors (NRC 1994). Thus, metal fuel does not have any clear safety advantages in this regard relative to other types of fuel.

In any event, the negative reactivity feedbacks from fuel expansion are not as fast-acting and therefore are less dependable than Doppler feedback. Shutdown mechanisms from fuel expansion “are somewhat delayed because of the inertia that must be overcome” (Lewis 1977). Such delays are problematic because fast reactors can experience “extremely rapid rates of power increase” if the system becomes “even slightly” supercritical (Lewis 1977).

MODERATOR AND COOLANT TEMPERATURE COEFFICIENTS

For LWRs, the moderator and the coolant are the same: ordinary water. If the coolant water heats up, it will expand and become less dense. This reduces the ability of the water, acting as a moderator, to slow down the neutrons as required for a thermal reactor—generally leading to a reduction in the reactor’s power output. Therefore, the temperature coefficient of reactivity for the coolant is mostly negative, although it can be slightly positive for some operating conditions (which are strictly limited by the regulator).

For fast reactors, especially those that use plutonium fuel, the opposite is often true. Fast reactors rely on fast neutrons to maintain a chain reaction and use coolants, like liquid sodium, that do not have a significant moderating effect. Nevertheless, neutrons do lose some energy when they collide with the coolant nuclei. As the coolant gets hotter and becomes less dense, the neutrons collide less frequently with the sodium coolant, and the population of neutrons becomes slightly more energetic, which increases the probability that the plutonium-239 fuel will fission if struck by a neutron and increases the number of neutrons released per fission. Both these effects can increase the reactor power. However, at the same time, the less dense coolant allows more neutrons to leak out of the core, which could reduce the power. Thus, there are two competing effects. The sign of the coolant temperature coefficient could be either positive or negative, depending on which effect is dominant in a specific core.

VOID COEFFICIENT

As the power level and temperature of an LWR increases, the water will eventually reach the boiling point and form steam (or if the coolant is already boiling—as in boiling-water

reactors—the steam bubbles, or voids, will expand). The formation of voids decreases the density of the moderator at a faster rate than an increase in temperature without a phase change, and will also reduce reactivity.

For LWRs, where water is both the coolant and moderator, the formation of steam bubbles will result in less moderation of the neutrons and hence a reduction in the power output. Thus, these reactors generally have negative void coefficients. For sodium-cooled fast reactors, the void feedback effect is quite different. Void formation in the sodium will significantly and rapidly reduce its density. As is the case with the coolant temperature coefficient, the void coefficient could be either positive or negative, depending on the properties of the reactor system and where the voids occur.

The sodium void coefficient depends strongly on the size and shape of the reactor core. For a fixed shape, larger cores will leak fewer neutrons than smaller cores (because the neutrons have to travel farther through the core before they escape, increasing the chance that they will strike a nucleus). Therefore, the void coefficient tends to become more positive as the core size increases. And for a fixed core volume, the rate of neutron leakage depends on the shape. The more surface area is available for neutrons to escape, the higher the leakage rate. The 471 MWth PRISM reactor has a positive void coefficient (NRC 1994, Petti et al. 2017). The 300 MWth VTR is reported to have an overall negative sodium void coefficient because of its size and shape, but has a locally positive coefficient near the center of the core where neutron leakage is less probable (Heidet 2019).

Because of the safety risks from positive power feedback, some engineers have tried to design large fast reactor cores with negative or very small positive void coefficients. This has turned out to be very challenging. One approach to reducing the void coefficient is to make the core leakier by changing the shape of the reactor so that it has a relatively high surface area to volume ratio (such as a pancake does). However, this is hard to implement in practice and can decrease reactor performance and safety in other ways. For instance, a fast reactor’s capability to breed—one of the major advantages cited by advocates—depends on how efficiently the reactor uses neutrons to convert U-238 to plutonium. Fast reactors designed to leak a lot of neutrons have worse breeding performance.

As a result, some fast reactor designers have concluded that it is not necessary to eliminate the positive void coefficient but only to add design features to mitigate its impact; others

CONTINUED ON P. 63

BOX 6 CONTINUED

are not ready to accept this compromise. One of the original goals of the now-cancelled ASTRID project in France was to design a large (600 MWe) oxide-fueled fast reactor with an overall negative sodium void coefficient. However, this resulted in what one report called a “peculiar” core design that is very complex to analyze and that is still susceptible to positive feedback effects if sodium voids form in only one part of the core (Nuria et al. 2017). It is unclear whether it will be possible to design practical fast reactors with a negative sodium void coefficient.

RADIAL EXPANSION COEFFICIENT

Another source of negative reactivity with increasing temperature comes from the expansion of the core in the radial direction—that is, perpendicular to the axis of the fuel assemblies. This occurs primarily from the expansion of core structures, such as the metal grids that hold the fuel assemblies in place. In contrast to the fuel temperature coefficient, this is not an inherent property of the fuel, and it is not a prompt (nearly instantaneous) response, so it should not be considered as reliable a feedback mechanism as the fuel temperature coefficient.

continued from p. 60

response to changes in the state of the reactor, such as temperature and pressure. While LWRs generally have negative power feedback and meet the NRC’s General Design Criterion 11, fast reactors generally do not meet this criterion and have power instabilities that can have serious implications for safety.

PASSIVE SAFETY CLAIMS AND THE EBR-II 1986 SAFETY TESTS

As discussed above, proponents of metal-fueled sodium-cooled fast reactors claim that these reactors are passively safe and would shut themselves down with no operator intervention if the fuel were to overheat. They frequently point to a series of safety tests carried out by Argonne National Laboratory in 1986 at the 62.5 MWth (20 MWe) Experimental Breeder Reactor-II (EBR-II). The most severe of these tests, on April 3, 1986, known as a “loss of flow without scram,” is often cited as a conclusive demonstration that fast reactors are passively safe. In the 2013 documentary

Pandora’s Promise, a former EBR-II nuclear engineer described the experiment as “almost a direct parallel to what happened at Fukushima” and claimed that “the reactor quietly shut itself down” (Stone 2013).

In a loss of flow accident, a fast reactor’s primary coolant pumps stop operating, coolant flow through the core is greatly reduced and the fuel temperature increases. The reactor has a protection system that normally would shut it down automatically by triggering a “scram,” or the rapid insertion of control rods. However, there is a small chance that the scram will not work. For the EBR-II safety test on April 3, 1986, operators simulated a loss of flow without scram to observe whether the reactor would reach a stable state or continue to overheat. Operators brought the EBR-II to 100 percent power and then switched off coolant pumps to simulate the impact of a total loss of alternating current power (a station black-out, similar to what occurred at Fukushima in March 2011). Unlike Fukushima, however, the EBR-II was not scrammed to stop the chain reaction, and it continued to generate power. As expected, the reactor temperature rose rapidly. But the overall system exhibited a negative reactivity feedback effect in response to the increase in temperature, and the temperature then decreased until the reactor reached a stable state at low power (although it did not completely shut down). At this low power level, natural forces such as convection removed heat quickly enough to allow the temperature to stabilize. The reactor fuel remained intact because the temperature stopped rising before the fuel heated up to its damage point.

In a second type of test, operators simulated a “loss of heat sink without scram,” shutting off the secondary coolant system so that there was no way to remove heat from the reactor. This caused the primary coolant to heat up as the reactor continued to generate power. In this test as well (though a less severe challenge than the loss of flow without scram), the resulting temperature increase caused negative reactivity feedback that safely shut down the reactor without causing fuel damage.

THE LIMITATIONS OF THE SAFETY TESTS

These test results appear impressive at first glance, but there is less here than meets the eye. For example, in one test, the reactor operated for only a couple of hours before the test in order to limit the decay heat after scram (IAEA 2017). In others, the fission chain reaction never actually stopped. And although the reactor fuel was not damaged due to overheating in the loss of flow without scram tests, in three tests the final fuel temperature did exceed the safety limit that operators had established—including the one on April 3, 1986.

Most significantly, the entire series of tests was highly scripted and conducted under very carefully controlled conditions to minimize the potential for failure; the EBR-II underwent “a number of hardware and software changes” to prepare the facility for the tests (Messick et al. 1987). While it is reasonable—or even essential—to carry out such preparations when doing safety tests, it limits the tests’ relevance to real-world accident scenarios. At best, the tests demonstrated certain safety features were functional under certain conditions, but they did not simulate the entire range of plausible severe accident conditions.

Thus, the tests did not prove that the reactor “cannot melt down,” as some claimed in Pandora’s Promise. On the contrary, as discussed below, fast reactors are vulnerable to a number of different accident initiators that could result in core damage and release of radioactivity into the environment. Below are some of the parameters of the EBR-II safety tests that had significant impacts on the outcomes and prevent the test results from being applicable to all fast reactors in all accident scenarios of concern.

Coastdown Time

A critical factor for the success of the EBR-II “loss of flow without scram” tests was the relatively long period of time that the primary system coolant pumps took to coastdown (gradually slow down after the power was cut off). This is because most of the feedback mechanisms that work to reduce a fast reactor’s power need time to take effect. If the coolant pumps shut down immediately, then the sodium temperature could increase to its boiling point before the negative reactivity feedback would have a chance to kick in and reduce the reactor power. Therefore, fast reactor coolant pumps must be designed with a sufficiently long coastdown period to allow the reactor to withstand a loss of flow without scram accident without fuel damage. The EBR-II loss of flow without scram tests extended the coastdown time artificially, and thus were not fully representative of real accidents.

The EBR-II used two conventional motor-driven centrifugal pumps to circulate the primary coolant, as well as an auxiliary electromagnetic pump with a battery backup to supplement natural convection cooling during shutdown. For motor-driven centrifugal pumps, which drive fluid motion by rotating, coastdown occurs naturally due to the rotational inertia of the pump. However, the natural coastdown period of a pump may be too short, requiring artificial means to lengthen it. And electromagnetic pumps—which induce electromagnetic fields to drive metal coolant flow and have no moving parts—require auxiliary mechanical flywheels

to simulate coastdown if the pumps stop operating. These artificial mechanisms do not fully compensate for pumps with short natural coastdown periods because they could be vulnerable to failure during an accident.

The scientists at Argonne National Laboratory knew that pump coastdown was “critically important in determining the peak transient temperatures” (Planchon et al. 1987). However, the 50-second coastdown period of the centrifugal pumps was too short. So to prepare for the 1986 loss-of-flow tests, the scientists introduced electronic controls and other modifications to artificially extend the coastdown time to as long as 600 seconds (Planchon et al. 1987). Given that rotating pumps can seize—suddenly stop running—a sufficiently long coastdown time is not guaranteed and should not be considered an intrinsic passive safety feature.

In addition to the artificial lengthening of the primary pump’s coastdown time, other parameters that were varied included the operating power of the reactor and the state of the auxiliary coolant pump (on, off, or on battery backup) (Planchon et al. 1988).

However, none of the tests included the most challenging but still plausible combination of conditions: 100 percent power, natural coastdown time for the primary pumps, and no auxiliary pump. In short, the tests did not provide information about how the reactor might respond to complex, real-world accidents that evolved in unexpected ways, as was the case at Fukushima. For example, the April 3, 1986, test was initiated at full power, but the coastdown time of the primary pump was extended to 95 seconds and the auxiliary pump was allowed to operate on battery power. Thus, the test differed in an important way from the Fukushima accident, when nearly all battery power supplies, as well as the electric distribution systems, were lost due to the flooding from the tsunami. Moreover, at Fukushima, fuel melting did not begin until several hours after all power (alternating current as well as battery) was lost. But in a fast reactor, fuel melting can begin within seconds after a total loss of power causing the failure of both primary and auxiliary pumps.

Sodium Void Coefficient

Another major reason why the EBR-II safety tests were not representative of all fast reactors is because the reactor had a negative sodium void coefficient (Chang 1992), which would not normally be the case for full-scale power reactors. As discussed earlier, larger fast reactors such as PRISM (165 to 311 MWe) will typically have a positive sodium void coefficient. As the size of the reactor core increases, the fraction of neutrons that leak from the core decreases. Also, plutonium-fueled reactors will have larger positive void coefficients

than uranium-fueled ones, with potentially more severe consequences. In a reactor with a positive sodium void coefficient, if the sodium coolant heats up and starts to boil, the total reactivity and power will sharply increase, heating the coolant even more and resulting in a positive feedback loop. This is an unstable condition that could result in a power increase so rapid that it would be impossible to control—as occurred during the 1986 Chernobyl accident.

The EBR-II's void coefficient was negative because it was a small reactor and was fueled with HEU instead of plutonium. The small size of the core enhanced the neutron leakage, and the use of U-235 instead of plutonium resulted in a smaller increase in fission rate at high neutron speeds. The negative void coefficient worked in concert with other phenomena contributing to negative feedback, such as expansion of the core. Indeed, the negative void reactivity was the largest contributor to the overall negative temperature feedback (IAEA 2017). If the void coefficient of the EBR-II had been positive, the other feedback effects might not have been strong enough to offset its impact on reactivity, and the loss of flow without scram tests would have had less benign outcomes. Therefore, the EBR-II tests do not demonstrate that all metal-fueled fast reactors are passively safe.

Notably, the NRC itself has questioned the use of the term “passively safe” to describe reactors with a positive void coefficient such as PRISM (NRC 1994):

The positive sodium void worth is a concern in the passive safety argument. Because of it, one must qualify any characterization of the PRISM response as “passively safe” by pointing out that this is conditional on the sodium remaining below the boiling temperature. Should sodium boiling begin on a core-wide basis under failure-to-scram conditions, the reactor would be likely to experience a severe power excursion.

In other words, if the liquid sodium boils, the reactor power would continue to rapidly increase, overwhelming the passive safety features, and a severe core damage accident could result.

How likely is it that the sodium would boil during an accident? Fast reactor developers argue that such an event would be extremely unlikely because there is a significant margin between the normal operating temperature of the reactor (around 500°C) and the sodium boiling point (around 900°C). Nevertheless, the likelihood of a rapidly developing sodium boiling event is design-dependent, highly uncertain, and not so easily dismissed.

There is very little information about the temperature limits of metal fast reactor fuel and how much time would

be available before fuel damage would occur if cooling were lost (NRC 1994). However, a loss of flow without scram would likely result in a devastating accident if all of the primary coolant pumps were to seize abruptly. In that case, the NRC's analysis found that large-scale sodium boiling would begin after about 14 seconds, leading to a power increase after 25 seconds. By 26 seconds, the power level would have increased by a factor of three, and the temperature at the centers of the fuel pins would have exceeded 1300°C, which is greater than their melting point. The NRC terminated the calculation at 26 seconds because there was little doubt where things were headed after that. The NRC report dryly states that “assuming that the prediction of the sodium flow rate through the core . . . is correct, this is clearly an event that must be avoided” (NRC 1994).

Another type of fast reactor accident known as the unprotected transient overpower event, in which a control rod is ejected and the reactor fails to shut down, could also be very severe. An Argonne National Laboratory analysis found that such an event at a relatively small SFR (380 MWe, similar to the Natrium) could cause large-scale fuel melting within 10 seconds, and dangerously high radiation doses to the off-site public (hundreds of rem at a 200-meter site boundary) (Grabaskas et al. 2016). These doses are not lower than those that could result from a core-melt accident at a large LWR.

THE HYPOTHETICAL CORE DISRUPTIVE ACCIDENT

It is commonly said that nuclear reactors cannot explode like nuclear weapons. While this is essentially true for LWRs, it is not the case for other types of reactors, such as Chernobyl-type thermal reactors or fast reactors. In the event of a severe accident in which nuclear fuel overheats, the fuel elements may melt and fuse together into a dense mass. The consequences of this compaction will differ in LWRs and in fast reactors.

Because LWRs require a moderator (water) to be intermingled with the fuel to achieve criticality and produce power, if the fuel becomes more compact and the moderator (water) is expelled from the core, there is a greater chance that neutrons will be absorbed in the fuel or escape before they are slowed down enough to cause fission and generate more neutrons. This reduces the reactivity of the reactor and increases its stability.

However, in a fast reactor, a moderator is not needed to achieve criticality. If the fuel rapidly becomes denser, then there is a smaller chance that neutrons will leave the fast reactor system without causing fission. This *increases* the reactivity and power of the reactor, not unlike the mechanism

by which nuclear fission weapons explode. In a nuclear fission weapon, plutonium or another fissionable material is rapidly assembled or compressed to a highly supercritical state. The yield of the weapon depends on the degree of compaction and how long it remains in a supercritical state before blowing itself apart. Although the degree and time scale of the compaction in a fast reactor would not lead to conditions nearly as destructive as most nuclear weapons, a small explosion is possible, with the potential to cause a catastrophic dispersal of highly radioactive fuel into the environment.

Hans Bethe, the Nobel laureate in physics who headed the theoretical division for the Manhattan Project, and a collaborator, J. H. Tait, were quick to recognize this risk and developed a back-of-the-envelope method to estimate the potential explosive energy release. When applied to metal-fueled fast reactors such as EBR-II and Fermi 1, the Bethe-Tait method revealed that the resulting explosion could be comparable to a detonation of several hundred pounds of high explosive (Lewis 1977). Although not as large as typical truck bombs today, such an explosion in the core of a nuclear reactor could breach the reactor vessel and containment. The resulting radiological release would not be primarily composed of radionuclides generated by the fissions during the explosion (as in the case of a nuclear weapon detonation), but rather those that accumulated during operation of the reactor—generally a much larger quantity of long-lived fission products.

This type of analysis, which was subsequently refined by others, also revealed that the size of the explosive energy detonation increases with the volume of the reactor core and could be significantly greater for commercial-scale reactors than for the relatively small EBR-II core (Lewis 1977). This type of scenario was christened the hypothetical core disruptive accident or HCDA. The modifier “hypothetical” was originally added to emphasize that the initial analyses simply assumed the core could become compacted but did not postulate how that could actually occur. However, researchers have identified plausible fast reactor accident sequences that could result in core meltdown and rapid reactivity increases (Tentner et al. 2010). Moreover, after the meltdowns at Three Mile Island in 1979, Chernobyl in 1986, and Fukushima in 2011, the occurrence of severe accidents is no longer hypothetical—and the “H” in HCDA is unnecessary. Nevertheless, the term sometimes continues to be used.

Historically, the question of considering HCDAs in the design and licensing of fast reactors such as the Fast Flux Test Facility at the DOE’s Hanford site in the state of Washington and the proposed (but never built) Clinch River Breeder Reactor in Tennessee was very controversial.

Reactor developers feared that it would prove too costly if regulators required that reactor structures be engineered to withstand HCDAs. Ultimately, the NRC decided in the 1970s that HCDAs could be excluded from the design basis for Clinch River if their probabilities were shown to be sufficiently low, which the applicant claimed was the case (Flanagan, Fanning, and Sofu 2015).

Nevertheless, the events are complex, the uncertainties large, and the potential consequences catastrophic. An [H]CDA is a very real and serious risk that must be considered in evaluating the prospects for fast reactors. The Generation IV International Forum, in a recent assessment of sodium-cooled fast reactor safety, concluded that “the possibility to robustly mitigate consequences of [a] whole core accident has to be investigated,” including the need for “a robust confinement capability” for radioactive material releases (Ruggieri et al. 2017).

Sustainability and Proliferation/ Terrorism Risk

The above discussion shows that fast reactors have troublesome features that may render them less safe than LWRs. But some of these very characteristics do offer the potential for increased sustainability compared to LWRs. As discussed in chapter 3, fast reactors are capable (in theory) of significantly improving uranium utilization by breeding plutonium, or significantly increasing the capacity of geologic repositories by more effectively fissioning (or “burning”) the long-lived TRU in spent nuclear fuel.

However, there are two very large caveats. As discussed in chapter 3, it is completely impractical, if not impossible, to achieve either of these sustainability objectives within a realistic timeframe. Depending on the technologies employed, it could take centuries or even millennia for fast burner reactors to recycle a significant fraction of the TRU contained in spent fuel, or for fast breeder reactors to utilize a significant fraction of the depleted uranium stockpile.

And as discussed in chapter 4, reprocessing and recycle of plutonium and other TRU greatly increases the likelihood that nations or terrorists seeking nuclear weapon-usable material will succeed. The security and safeguards measures needed to mitigate these risks are costly and cumbersome, yet only have limited effectiveness.

Sustainability of Once-Through Fast Reactors

A number of vendors are developing fast reactors that would utilize high-assay low-enriched uranium (HALEU) fuel in a once-through cycle without reprocessing, at least for their initial operation. While these reactors, such as TerraPower’s

Natrium, would pose lower proliferation and nuclear terrorism risks than TRU-fueled fast reactors with reprocessing, they are not more sustainable than LWRs and thus fail to realize one of the main benefits of developing fast reactors. Indeed, such fast reactors are typically even less uranium-efficient than LWRs, as the below discussion illustrates.

Although TerraPower has provided few details publicly about the Natrium’s fuel cycle, enough information is available about it and similar reactor designs to estimate its uranium utilization efficiency. An 18.75%-enriched HALEU-fueled fast reactor core requires approximately 2.5 times more natural uranium per GWe than LWRs (Hoffman and Fei 2019). About 820 metric tons of natural uranium would be needed for the core of a typical 1,000 MWe LWR containing 89 metric tons of LEU with an enrichment of less than 5 percent. Thus about 2,050 metric tons of natural uranium would be needed for the core of a 1,000 MWe HALEU-fueled fast reactor. For the 345 MWe Natrium fast reactor, this corresponds to about 710 metric tons of natural uranium, which would be enriched to produce 17.6 metric tons of 18.75 percent HALEU. This value is consistent with the range of 15–20 tons per startup core that a Terra-Power representative specified at a 2020 DOE workshop (Gallacher 2020).

In addition to the natural uranium needed for the initial core fuel load, the annual natural uranium requirement for a HALEU-fueled fast reactor utilizing current fuel technology will also be greater than for an LWR of similar capacity. TerraPower has said that the Natrium will operate on a 18–24 month refueling cycle typical of LWRs. Also, GE-Hitachi has specified that the average discharge burnup of a HALEU-fueled PRISM reactor would be about 70,000 MWd/MTHM (Petti et al. 2017).¹⁵ If one assumes that the HALEU fuel will be irradiated for three 18-month cycles (4.5 years in total) and that the Natrium has a capacity factor of 85 percent, the reactor will require about 3.7 metric tons of HALEU per year, corresponding to a natural uranium requirement of 150 metric tons per year, or 500 metric tons/GWe-year. This is about 2.5 times the annual natural uranium requirement for an LWR.

HALEU-fueled once-through fast reactors also generate more long-lived radioactive waste than LWRs. The quantity of TRU discharged in the spent fuel per GWe-year would be over 500 kg per year, or more than twice the comparable value for an LWR (Hoffman and Fei 2019).

The benefits for sustainability of fast reactor-based fuel cycle systems with reprocessing are modest at best, but the proliferation and nuclear terrorism risks of such systems are profound. And while fast reactors operated with HALEU fuel on a once-through basis are less risky, they also are less uranium-efficient than LWRs. Thus it is likely that reactor developers

are pursuing once-through HALEU-fueled fast reactors only as “gateway” reactors to facilitate a transition to TRU-fueled reactors with reprocessing.

Time Scale and Costs

READINESS FOR COMMERCIAL DEMONSTRATION

Recall from chapter 1 that the commercialization of a new reactor technology involves several stages: research and development, an engineering demonstration using a reduced-scale prototype reactor, a performance demonstration using a scaled-up prototype, and a commercial demonstration using a full-scale reactor that would be the model for subsequent units. A reactor that has had a successful performance demonstration is usually considered “high maturity.”

When estimating the time scale and cost to commercialize an NLWR, it is also necessary to consider the fuel cycle facilities necessary to support the reactor’s operation. These facilities may use experimental technologies that have not been demonstrated at commercial scale. Before the reactor can be commercially deployed, all the elements of the fuel cycle infrastructure—commercial-scale fuel production, transportation, and spent fuel management—must be available and reliable.

The DOE’s 2017 study of advanced demonstration and test reactor options identified sodium-cooled fast reactors as a high-maturity technology, estimating that it would cost \$4 billion and take 13 to 15 years to build and start up a commercial-scale demonstration reactor. The analysis concluded that that the 471 MWth (165 MWe) PRISM design is ready for a commercial demonstration, because prior fast reactor projects “had engineering and performance demonstration systems over three decades ago” (Petti et al. 2017). Consistent with this conclusion, the DOE has now chosen two PRISM-type designs—the Versatile Test Reactor and the Natrium commercial demonstration reactor—for near-term deployment.

However, the DOE has not made the case that PRISM-based fast reactors are mature enough to bypass the performance demonstration stage. First, prior fast reactor performance demonstrations were less than successful. The DOE report points to the Fermi-1 and Fast Flux Test Facility reactors as US examples and Phénix and Superphénix (France), Monju (Japan), and BN-600 (former Soviet Union) as international examples. But Fermi-1 and Monju both suffered major accidents. Phénix experienced operational anomalies that remain unexplained. Superphénix also never achieved full power and was unreliable. And the BN-600 experienced many sodium fires. The DOE report itself acknowledges that

“the track record of [sodium-cooled reactor] demonstration plants is mixed” (Petti et al. 2017).

Second, and even more to the point, those reactors differed in significant ways from the VTR and Natrium designs and thus are not directly relevant. The 20 MWe EBR-II was the experimental fast reactor most similar to PRISM. Could data from the EBR-II, a smaller-scale engineering demonstration, enable PRISM to leapfrog over the performance demonstration phase? According to the DOE’s definition, as discussed in chapter 1, performance demonstration is needed to establish that scale-up of the system works, gain operating experience to validate integral behavior of the system, and provide proof of performance. However, the results of the EBR-II demonstration cannot be readily extrapolated to larger PRISM-type reactors such as the VTR and the Natrium.

The 840 MWth Natrium would have a power rating 14 times that of the EBR-II—a significant difference with regard to many aspects of reactor operation. Indeed, GE-Hitachi originally recommended in 2016 that the DOE build a smaller, 471 MWe PRISM demonstration reactor rather than the larger, 840 MWth model (Petti et al. 2017). But the Natrium would be more than a mere EBR-II scale-up. Unlike the EBR-II, the sodium void coefficient for a metal-fueled fast reactor as large as the Natrium will be positive. Also, the DOE notes that the EBR-II cooling system “was effective for a reactor of that size” but “a larger reactor may require a different technology” (Petti et al. 2017). The cooling systems that work for very small reactors may well be inadequate for larger reactors that generate heat at a higher rate. For instance, PRISM designs includes a novel, additional passive decay heat removal system called the Reactor Vessel Auxiliary Cooling System, which the EBR-II did not have.

In addition, the design of PRISM’s primary coolant pumps would be different than those of the EBR-II, which could significantly impact safety analyses. As noted above, for ordinary motor-driven centrifugal pumps, coastdown occurs naturally as a consequence of the rotational inertia of the pump after power is cut off. The EBR-II had two primary centrifugal pumps and one auxiliary electromagnetic pump. However, PRISM designs use only electromagnetic pumps, which take advantage of the fact that the coolant is metallic. However, as discussed above, these pumps do not have moving parts and therefore have no intrinsic inertia to allow for coastdown. To compensate, the PRISM design includes synchronous motor-generator machines that are intended to simulate coastdown by providing power for a short time to the electromagnetic pumps in the event of a station blackout (NRC 1994). However, this arrangement is not as safe as a centrifugal pump. For instance, if the power connections

between the motor-generator machines and the electromagnetic pumps are disabled—in the event of a severe flood, for instance—then there would be no coastdown effect, and a rapid power excursion and core meltdown could result.

Also, the equipment needed to provide high coolant flow rates following a pump shutdown is expensive, according to a recent study on the DOE VTR project (Sumner and Fanning 2020). As a result, DOE researchers are looking for ways to cut the VTR’s cost by reducing the post-trip flow rate, which would lower the safety margin as well (Sumner and Fanning 2020). Thus to reduce cost, commercial fast reactors based on the PRISM design may have lower post-shutdown flow rates than the EBR-II did, increasing the risk of an accident.

The NRC recognized in the 1990s that the EBR-II was not representative of the PRISM design and other fast reactors in its 1994 PRISM preapplication safety review:

... the fact that EBR-II is obviously quite different from the other cores decreases one’s confidence in extrapolating from the EBR-II test series. Analyses consistently indicate that the “passive shutdown” will work as designed in the PRISM, but a series of safety tests using a prototype reactor is needed for confirmation (NRC 1994).

Another important difference is that neither the EBR-II nor other fast reactor demonstrations used a fuel similar to the uranium-plutonium-zirconium metal alloy that the VTR will use. Fermi-1 used a metallic fuel alloy of highly enriched uranium (25.6 percent U-235) and molybdenum, which was bonded to the cladding in a different way than PRISM fuel. The FFTF and other reactors used mixed plutonium-uranium oxide (MOX) fuel.

Almost all the PRISM-type driver fuel used by the EBR-II was an HEU-containing metal alloy. Less than 1 percent of the fuel tested in the reactor was composed of a plutonium-uranium-zirconium alloy, which would be the fuel of choice not only for the VTR but also for future fast breeder reactors. And the EBR-II did not test metal fuel with TRU other than plutonium, which would be included in the fuel for a TRU burner fast reactor. Such differences—highly enriched uranium versus plutonium, metal versus oxide—could have significant impacts on reactor operation, safety, and performance.

The NRC also flagged the fuel issue in its 1994 assessment:

The PRISM fuel system, U-Pu-Zr fuel clad with HT9, is a new concept. Many of the basic design principles have been developed from EBR-II metal-fuel experience. However, because of differences in material, geometry, and exposure conditions, this experience must be extrap-

olated to the PRISM design through the use of analytical tools that characterize the operational history and transient responses of the fuel system. Experimental data must be obtained both to support the model development efforts and to verify the integrated computer codes (NRC 1994).

The lack of experience with plutonium fuel will not be an issue for the HALEU-fueled Natrium. However, there is disagreement about the maturity level of the proposed PRISM fuel—even without plutonium—in light of past EBR-II experience. The DOE’s evaluation team judged that it was significantly less mature than the GEH team did. The DOE noted that “no attempt was made to reconcile this [its own] independent assessment with those of the individual team assessments” (Petti et al. 2017).

At the time of the NRC’s 1994 PRISM preapplication review, the DOE planned to conduct additional research and development for metal fuel qualification at the EBR-II and also build a more representative prototype of a PRISM module to conduct full-scale performance and safety testing, which would have addressed the NRC’s concerns. The DOE had laid out a 20-year PRISM technology development schedule to support an NRC design certification.

However, the EBR-II was shut down in 1994, the DOE never built the prototype, and the required fuel qualification and safety testing was never carried out. An additional fuel transient testing program to support severe accident analyses was supposed to take place at the Transient Reactor Test (TREAT) Facility, a test reactor at the INL, but was never carried out because TREAT was also closed in 1994. The DOE recently restarted the TREAT reactor to conduct transient fuel testing, but that effort is only in its beginning stages, and it will take many years to accumulate enough data to make a strong safety case for fast reactor licensing.

The decision to bypass prototype testing raises questions about how the safety case for the VTR and Natrium will be adequately demonstrated and how their fuels will be qualified. Both reactor designs are very different from EBR-II, as discussed above. The DOE’s current approach, to rely on EBR-II performance and fuel data, would suffer from many of the same issues raised by the NRC for PRISM licensing.

If the proposed VTR project goes forward, it could perform some of the fuel qualification activities for commercial fast reactors that would have been done by the EBR-II in the 1990s. However, the VTR is not likely to be operational before the late 2020s, and even then it will likely require several years of commissioning activities before it can begin sustained operation. Thus fuel qualification programs at the VTR could take well into the 2040s, taking into account irradiation time,

post-irradiation examinations, and additional safety testing. And if the VTR is used to qualify its own fuel, there will be a bootstrapping problem that could raise safety concerns.

It is unclear why the DOE has now changed its position from the one it held in the 1990s, and now maintains that construction and testing of a PRISM prototype will not be necessary prior to licensing a commercial demonstration reactor (or, by the same logic, a large test reactor). An open question is whether the NRC which has regulatory authority over the Natrium demonstration reactor but not the VTR, will also agree that it can proceed with licensing the Natrium without requiring prototype testing.

To recap, compared to the EBR-II, PRISM-based designs such as the VTR and Natrium are many times larger, use only electromagnetic coolant pumps, may use fuel containing plutonium and possibly other TRU, and will have positive sodium void coefficients. All of these factors tend to make accidents such as the loss of flow without scram more severe. Therefore, the relevance of the EBR-II safety tests to commercial-scale PRISM systems is highly questionable. There is little evidence to support the DOE’s assertion that the PRISM design is ready for deployment as either a commercial demonstration or as a test reactor without first conducting a performance demonstration for safety testing.

If the 20-year schedule that the DOE proposed in the 1990s for PRISM development were followed today, with credit for preliminary design activities that began around 2018, licensing of the first commercial unit would not take place until the late 2030s. Assuming that there were commercial orders at that time, the first units would not likely be operational until around 2050.

FUEL CYCLE READINESS

In order for commercial PRISM-type units such as the Natrium to be available by the 2030s, all the fuel cycle facilities needed to support reactor operation would also need to be available. As noted in chapter 3, two of the primary justifications for fast reactor development are to reduce the generation of long-lived radioactive waste and to use uranium more efficiently. As discussed above, the once-through, HALEU-fueled Natrium will not achieve either of these objectives. To do so, the Natrium and most other fast reactor concepts would require reprocessing of their spent fuel to recover uranium, plutonium, and possibly other TRU for use in fresh fuel—using technologies that themselves require intensive development. For example, the full PRISM fuel cycle would require facilities for pyroprocessing the reactor’s metal-based spent fuel to extract plutonium and other TRU, and facilities for fabricating fresh fuel from the separated materials.

Moreover, the current PRISM metal fuel design, which is also the basis for the VTR, the Natrium, and Oklo's Aurora microreactor, contains a bond of metallic sodium between the fuel and the metal cladding to provide for good heat transfer. The DOE has long argued that such fuel, if exposed to a high burnup cannot be directly disposed of in a repository but must be processed to remove the bond sodium. Unless the DOE changes its position—and unless new fuel designs without bond sodium are qualified for safe use in the future—it will require processing of the spent fuel from all three of these fast reactor projects.

The process that the DOE has used to date for reprocessing metallic fast reactor spent fuel, pyroprocessing, has not been demonstrated at commercial scale, and its only significant operating experience at a reduced scale arguably has not been successful (see Box 7, p. 71). Since 1996, US researchers at the INL have been struggling to pyroprocess 26 metric tons of metallic spent fuel from the shutdown EBR-II and FFTF fast reactors. Of this amount, about 3.2 metric tons consists of HEU-based driver fuel, which is being down-blended with depleted uranium to produce HALEU with an enrichment just under 20 percent.

In 2000, the DOE estimated that the project would be completed by 2010, but as of December 2020, only about 20 percent of the spent fuel had been processed. At that rate, it appears likely that several more decades will be needed to finish the job (see Box 7, p. 71). About half of the driver fuel was pyroprocessed by 2020, at an average rate of about 85 kilograms per year. To put that in perspective, the VTR, which would discharge about 1.8 metric tons of spent driver fuel per year, would require a pyroprocessing annual throughput 20 times higher than this average rate. The 345 MWe Natrium would discharge about twice as much spent fuel per year as the VTR, or more than 40 times the current pyro-processing rate.

The DOE has plans to increase the operating time of the pyroprocessing facility by running it 24 hours a day, 7 days a week, which would only increase the throughput by a factor of about three (INL 2020). But it is not just a scale-up issue: The process at the INL is generating nuclear waste streams that cannot be effectively managed and do not have an established disposition path (see Box 7, p. 71). More fundamental development work is also needed at the engineering scale to address this waste problem.

The difficulty of scaling up the pyroprocessing process and other considerations, including the proliferation risks of separating plutonium, have led the DOE to recently propose a new and different process called “melt-distill-dilute” (also called “melt-distill-package”) for treating the VTR spent fuel (Crawford 2020). This process would involve melting the

spent fuel pins and then heating the melt until the sodium is driven off. The process would also vaporize volatile fission products, which would have to be trapped and stored. The plutonium and other TRU would not be separated from uranium, and would be diluted by fuel structural materials to below 10 weight-percent to meet the DOE's safeguards requirements. This concept still requires developmental work (Crawford 2020).

The DOE has not provided similar information about its proposed disposition path for the even greater quantity of sodium-bonded spent fuel that would be discharged by the Natrium annually. However, the melt-distill-package process will obviously not be suitable if the intention is to improve the unfavorable sustainability characteristics of the design—albeit modestly—by reprocessing and recycling the plutonium and other TRU in the spent fuel.

FUEL CYCLE FACILITY DEVELOPMENT

What additional time and resources would be needed to develop and build commercial-scale pyroprocessing and fuel fabrication demonstration facilities for sodium-cooled fast reactors?

A 2014 DOE study estimated that it would cost \$12 billion to \$35 billion to achieve a first-of-a-kind commercial demonstration of a relatively high-maturity advanced fuel cycle, of which only one component facility (e.g., the reprocessing plant) would require engineering-scale demonstration (Wigeland et al. 2014). For less mature technologies, which would need to demonstrate several components of a fuel cycle at engineering scale, the study estimated costs of \$35 billion to \$75 billion to reach the same stage. The study estimates that it would cost hundreds of billions of dollars to transition to a new fuel cycle in the United States.

The 2014 DOE study did not estimate the total time it would take to achieve commercial demonstration of a fast reactor fuel cycle. But a schedule presented by a DOE official at a conference of the American Nuclear Society in 2017 showed that it was estimated to take 45 years from beginning the design of an LWR fuel reprocessing plant (which would be needed to produce the initial plutonium for the fast reactor) to beginning the operation of a full-scale fast reactor reprocessing plant (Paviet 2017).

Another issue that will affect cost is the need for HALEU by many fast reactor designs, as discussed in chapter 4. Uranium enrichment plants would have to be built or reconfigured to supply this material, and downstream conversion and fuel fabrication plants would have to be modified to handle the criticality risks of such materials. Security would

continued on p. 72

The Pyroprocessing Files

The Union of Concerned Scientists has reviewed documents that it obtained through the Freedom of Information Act (FOIA) that highlight the failure of the DOE's decades-long effort to chemically process 26 metric tons of sodium-bonded metallic spent fuel from the shutdown EBR-II reactor (Lyman 2017). Ostensibly, the purpose of the program was to convert the waste to forms that would be safer for disposal in a geologic repository. A secondary goal was to demonstrate the viability of a new type of pyroprocessing process (see chapter 4). Instead, it has demonstrated the numerous shortcomings of this technology. Pyroprocessing, like other reprocessing technologies, takes one form of nuclear waste and converts it into multiple different types of nuclear waste, each presenting new challenges for disposal.

Pyroprocessing is a form of spent fuel reprocessing that dissolves metal-based spent fuel in a molten salt bath (as distinguished from conventional reprocessing, which dissolves spent fuel in water-based acid solutions). Understandably, given all of the technology's problems, the DOE has been reluctant to release public information on this program, which has largely operated under the radar since 2000.

DOE PYROPROCESSING ACTIVITIES

The DOE initiated the pyroprocessing program for EBR-II spent fuel in the mid-1990s as a consolation prize to Argonne-West National Laboratory (now part of present-day INL) after it cancelled the Integral Fast Reactor project. The idea was to connect the EBR-II to an adjacent pyroprocessing facility, which would extract plutonium, uranium, and other elements from the reactor's spent fuel and fabricate them into new reactor fuel. In theory, this could be a system that could convert its nuclear waste into usable fuel on site and thus be largely self-contained. Pyroprocessing was billed as a simpler, cheaper, and more compact alternative to the conventional aqueous reprocessing plants that have been operated in France, the United Kingdom, Japan, and other countries.

Although the DOE shut down the EBR-II in 1994 (the reactor part of its Integral Fast Reactor Program), it allowed work at the pyroprocessing facility to proceed. It justified doing so by asserting that the leftover spent fuel from the EBR-II could not be directly disposed of in the planned Yucca Mountain repository because of the potential for metallic sodium in the fuel (used to bond the fuel to the cladding) to react violently with water and air.

Pyroprocessing would separate the sodium from other spent fuel constituents and neutralize it. The DOE decided

in 2000 to use pyroprocessing for the entire inventory of leftover EBR-II spent fuel—both driver and blanket fuel—even though it acknowledged that there were simpler methods to remove the sodium from the lightly irradiated blanket fuel, which constituted nearly 90 percent of the inventory.

However, as the FOIA documents reveal in detail, the pyroprocessing technology simply has not worked well and has fallen far short of initial predictions. Although the DOE initially claimed that the entire inventory would be processed by 2007, as of the end of 2020, only about 20 percent of the roughly 26 metric tons of spent fuel had been processed. More than \$210 million had been spent, at an average cost of around \$50,000 per kilogram of fuel treated, compared to the original estimate of less than \$18,000 per kilogram. Since 2016, only driver fuel has been pyroprocessed, at a rate of about 100 kilograms per year and an annual cost of \$8 million (INL 2020). This corresponds to a cost of about \$80,000 per kilogram. At this rate, it would take until the end of the century to complete pyroprocessing of the entire inventory, at an additional cost of more than \$1 billion.

But even that assumes, unrealistically, that the equipment will continue to be usable for this extended time period. Moreover, there is a significant fraction of spent fuel in storage that has degraded and may not be suitable for pyroprocessing in any event. The long time to completion is problematic because the DOE has had an agreement with the state of Idaho since 1995 to remove all spent fuel from the state by the year 2035. The FOIA documents reveal that the DOE was well aware that it was not on track to comply with this obligation. In 2019, Idaho and the DOE reached a supplemental agreement with additional conditions, including a requirement that the DOE complete pyroprocessing of the EBR-II driver fuel by 2028. In order to accomplish this, the facility will have to be ramped up to 24 hour-per-day, 7 day-per-week operations by 2024, which is not realistic given the age of the facility and its previous operating record (INL 2020).

GENERATING AND ATTEMPTING TO MANAGE MULTIPLE WASTE STREAMS

What exactly is this pyroprocessing campaign accomplishing? Instead of making management and disposal of the spent fuel simpler and safer, it has created an even bigger mess. Pyroprocessing separates the spent fuel into three waste streams. The first is a cast metal ingot called the “spent fuel treatment product.” Some of this material is HALEU obtained by

CONTINUED ON P. 71

downblending highly enriched uranium recovered from pyroprocessed driver fuel with natural uranium to reduce the U-235 concentration. But because this material contains unacceptably high levels of plutonium and other contaminants, it has a high radiation dose rate and until recently it was considered a waste product. The material has been accumulating and taking up valuable space at INL storage facilities, causing its own safety issues.

In 2020, the INL reached an agreement with Oklo, Inc., to provide it with HALEU for its proposed Aurora fast micro-reactor at the laboratory. However, in order to use HALEU previously produced from the down-blending of EBR-II driver fuel, the material will have to be purified and recast into smaller ingots to reduce the dose rate and make it acceptable for reuse. As discussed in chapter 4, Oklo will require about 3 metric tons of HALEU for its demonstration reactor. The cost of supplying this fuel will be between \$50 and \$100 million.

The second waste stream is the molten salt bath that is used to dissolve the spent fuel. Fission products and plutonium have accumulated in this salt for 20 years. Eventually it will have to be removed and safely disposed of. But for various reasons—including cost and a lack of available space for the necessary equipment—the INL is reconsidering the original plan to convert this waste into a stable ceramic waste form. Instead, it may just allow it to cool until it hardens and then directly dispose of it in the Waste Isolation Pilot Plant in New Mexico.

The third waste stream consists of the leftover metal cladding tubes that encased the nuclear fuel and the metal plenums that extended above the fuel region. These tubes are

contaminated with fission products and sodium. The original plan was to convert these scraps into a stable, homogeneous high-level waste form. But the FOIA documents reveal that the DOE is considering redefining this material as transuranic or low-level waste so that without further processing it could be disposed of in the Waste Isolation Pilot Plant or a low-level radioactive waste disposal facility. Storage of the accumulating metal scrap material is also becoming an increasing burden at the INL.

A WASTE PROBLEM MAGNIFIED

Simply put, the DOE has spent hundreds of millions of dollars only to magnify, rather than simplify, the waste problem. This is especially troubling in light of other FOIA documents that indicate that the DOE never definitively concluded that the sodium-bonded spent fuel was unsafe to directly dispose of in the first place. But it insisted on pursuing pyroprocessing rather than conducting studies that might have shown that this costly, ineffective, and dangerous procedure was unnecessary.

Everyone with an interest in pyroprocessing should reassess their views given the real-world problems experienced in implementing the technology over the last 20 years at the INL. They should also note that the process needed to extract plutonium and other TRU to produce fresh fuel for fast reactors would be even more complex than the EBR-II pyroprocessing system and hence would require considerably more research and development. The technology is a long way from being demonstrated as a practical approach for recycling spent fuel for use in power generation.

continued from p. 70

have to be upgraded to Category II. However, HALEU supply is likely to be too scarce to meet demand for the foreseeable future. HALEU fuel will remain an expensive commodity unless significant new investments in supply and processing are made. In 2020, TerraPower announced a partnership with Centrus to expand its pilot centrifuge plant in Piketon, Ohio to produce HALEU for the Natrium demonstration reactor, but the facility would have to be quickly scaled up by a factor of three or more in order to produce the 15–20 MT of HALEU needed for the reactor's initial core by 2027. In the absence of sufficient domestic HALEU production, fast reactors in the United States may become dependent on

foreign producers, such as Russia and China, raising issues of reliability of supply.

Finally, fast reactors are significantly more expensive to build and operate than LWRs, and reprocessing and recycling spent fuel also increases cost, as discussed above. Thus, sodium-cooled fast reactors operating on a closed fuel cycle will likely generate more expensive electricity than LWRs on a once-through cycle until uranium becomes so scarce that its price increases to several times its current value. (And at that point, extraction of seawater uranium, a virtually inexhaustible resource, could be an economically competitive and more attractive alternative to reprocessing and plutonium recycling.)

In conclusion, the deployment of a fast reactor-based closed fuel cycle would likely decrease safety, would likely cost far more than LWRs on a once-through fuel cycle, and would make nuclear weapons materials more accessible to terrorists. And these reactors would neither solve the nuclear waste problem nor significantly reduce uranium use over reasonable time scales.

Fast reactors utilizing HALEU and operating on a once-through cycle with direct disposal of spent fuel would have many of the same safety risks but would be even less uranium-efficient than LWRs. The one exception, if the approach could work, would be a fast reactor that could operate in a once-through breed-and-burn mode (see chapter 8).

High-Temperature Gas–Cooled Reactors

The second NLWR concept identified by the 2017 DOE demonstration and test reactor study as being sufficiently mature to support a near-term commercial demonstration is the high-temperature gas-cooled reactor (HTGR), a thermal-neutron reactor (Petti et al. 2017). “High temperature” here is defined as an outlet temperature (where the coolant gas exits the reactor core) of up to 800°C.¹⁶ In comparison, pressurized-water reactors have a coolant temperature below 300°C. The higher coolant temperature makes this class of reactors about 20 to 33 percent more thermally efficient than an LWR and also would enable the reactor to provide high-temperature heat for industrial processes.

There are two primary types of HTGR. One is called a prismatic-block HTGR because the fuel elements are long blocks in the shape of a hexagonal prism. The second is a pebble-bed reactor, with spherical fuel elements. In contrast to prismatic-block HTGRs or LWRs, pebble-bed HTGRs are refueled continuously while the reactor is operating. Fuel pebbles are loaded at the top of the reactor core and circulate to the bottom, where they are removed. Depending on how long they have already been irradiated, pebbles are then either fed again into the reactor or stored as waste and replaced with fresh fuel.

Contemporary HTGR designs typically rely on passive means for emergency cooling, which limits their power to below about 300 megawatts-electric (MWe) in order to meet safety limits.

In October 2020, the DOE selected a pebble-bed HTGR design, the X-Energy Xe-100, as one of the two commercial demonstration plants to be built by 2027 under the Advanced Reactor Demonstration Program (ARDP).

The Technology

The HTGR uses graphite as a neutron-moderating material and helium gas as a coolant. Another of the HTGR’s distinctive characteristics is its fuel, which must be capable of withstanding far higher operating temperatures than LWR fuel. The current standard fuel, called TRISO (tristructural-isotropic), is composed of tiny spheres about one millimeter in diameter, each consisting of a kernel of fissile material (typically uranium oxide or uranium oxycarbide) surrounded by a porous graphite buffer layer, which is encapsulated in two spherical layers of pyrolytic carbon (a graphite-like material) with a silicon carbide layer sandwiched between them. Each layer serves a different purpose. The main objective of the layered fuel structure is to provide barriers to greatly inhibit fission product releases at the high temperatures this reactor would reach during normal operation and the even higher ones that could occur during design-basis accidents.

The TRISO fuel particles themselves are embedded in a matrix material in order to form fuel elements. There are two fuel element designs, corresponding to the two types of HTGRs. Prismatic fuel elements are fabricated by pressing TRISO fuel particles into a carbon matrix to form pellets called compacts, which are then inserted into holes drilled in prism-shaped graphite blocks. Pebble-bed fuel elements are fabricated by embedding 10,000 to 20,000 TRISO particles into graphite spheres 6 centimeters in diameter. The core of either type of reactor would contain billions of TRISO particles.

In principle, the special properties of TRISO fuel could lead to improved safety. The TRISO fuel coating can prevent

the release of most fission products up to a temperature of around 1600°C, whereas LWR fuel cladding begins to degrade and release some fission products at around 800°C. And the fissionable fuel particles are dispersed in a large volume of graphite, so an HTGR core has a lower power density and heats up more slowly than an LWR core if cooling is lost. In an LWR, the cladding failure temperature can be reached within minutes in the worst-case loss-of-coolant accident, while in an HTGR it could take tens of hours for the fuel to reach 1600°C.

However, there are caveats that make it difficult to assess whether HTGRs will be significantly safer overall than LWRs in practice. An HTGR must be designed to have a very low likelihood that its fuel temperature would exceed 1600°C during an accident, because the ability of TRISO particles to retain fission products decreases significantly if they heat up to higher temperatures. Also, TRISO fuel must be manufactured to very exacting specifications because the fuel will not perform as intended if it is produced incorrectly. This shifts part of the safety burden from the reactor to the fuel fabrication process. And a loss of coolant is not the only accident that could affect HTGRs. Other accident scenarios, as well as sabotage, could result in core damage and fission product release—some of which have not been thoroughly analyzed.

History and Current Status

The Department of Energy's (DOE's) judgment that HTGR technology is sufficiently mature for commercial demonstration is based on a considerable number of research, development, and demonstration projects over many decades. Between 2006 and 2016 alone, \$1.5 billion was spent by industry and the DOE on HTGR research and development (Kadak 2016).

HTGRs were first proposed in the late 1940s, and they initially generated a great deal of enthusiasm. In 1970, the Atomic Energy Commission even predicted that almost half of US nuclear capacity by the year 2000 would be made up of HTGRs (Shropshire and Herring 2004). General Atomics, a prominent HTGR developer, sold 10 large reactors to US utilities between 1971 and 1974 (Mcdowell et al. 2011). However, the HTGR revolution did not come to pass, in large part because the operating experience with demonstration reactors did not inspire confidence. Even the DOE concedes that “the track record of the early HTGRs is mixed” (Petti et al. 2017). Ultimately, all of the projects were cancelled due to “technological impasses and lack of competitiveness against light water reactors” (Shropshire and Herring 2004).

PAST DEMONSTRATION PROJECTS

Since the 1960s, there have been five HTGR engineering-scale demonstration reactor projects around the world, only one of which was in the United States: Peach Bottom 1 in Pennsylvania. There have also been two demonstrations of larger, commercial-scale plants: Fort St. Vrain in Colorado, and the THTR (Thorium Hochtemperatur Reaktor) in Germany (see Table 1 on p. 17).¹⁷

Engineering-scale prototypes such as the 115 megawatt-thermal (MWth) Peach Bottom 1, which operated from 1967 to 1974, performed relatively well, compared to the larger reactors. The 842 MWth (330 MWe) Fort St. Vrain reactor, a prismatic-block HTGR that was essentially a scaled-up version of Peach Bottom with technology improvements, operated from 1979 to 1989; during that time it experienced multiple technical problems and was highly unreliable. In Germany, the 300 MWe THTR, a pebble-bed HTGR, began generating electricity in 1985 and operated at full power for just two years before being shut down in 1989. It also experienced technical problems.

In recent decades, several other HTGR projects were initiated but failed to come to fruition. In the 2000s, the South African utility Eskom and the US utility Exelon pursued the development of a pebble-bed modular reactor in the United States. However, Exelon withdrew from the project in 2002, with the company's chief executive officer saying it was behind schedule and too speculative (Thomas 2008). Eskom continued development of the reactor in South Africa, ordering plant components and manufacturing fuel, but shut the project down in 2010 before the reactor was built.

The Energy Policy Act of 2005 authorized the Next Generation Nuclear Plant (NGNP) project, a commercial demonstration HTGR, at Idaho National Laboratory. However, that project was terminated after the private sector refused to commit to paying 50 percent of the research and development costs, as required by the Energy Policy Act (Kadak 2016). Subsequently, the DOE has continued to fund some HTGR-related research and development, including the Advanced Gas Reactor (AGR) fuel development and qualification program at Idaho National Laboratory and Oak Ridge National Laboratory.

A brighter spot for HTGRs may be China, where researchers have been developing the technology since the 1970s. In 1992, China decided to build a 10 MWth pebble-bed test reactor module. The HTR-10 uses a conventional steam cycle to produce electricity. The reactor reached first criticality in 2000 after five years of construction and achieved full-power operation in 2003 for 72 hours. Since then it has operated intermittently, primarily to conduct

experiments.¹⁸ Although it is difficult to find references documenting the HTR-10's operating performance, one report states that the reactor was down from 2007 to 2014 for maintenance, but does not provide further details (RAHP 2019). A project to couple the reactor to a gas turbine for electricity production by the end of 2005 was never carried out.

The HTR-10 is currently the only operating HTGR in the world. Japan's 30 MWth test HTGR, the HTTR (High-Temperature Engineering Test Reactor), has been shut down since the 2011 Fukushima accident. In June 2020, the Japanese Nuclear Regulation Authority authorized the reactor's owner, the Japan Atomic Energy Agency, to make changes to bring the HTTR in compliance with post-Fukushima regulatory standards, paving the way for its eventual restart by March 2021 (Ohno 2020).

NEW PROJECTS

Soon after the HTR-10 was first started up, China decided to build a much larger pebble-bed HTGR commercial demonstration plant, the HTR-PM, as well as a pilot fuel fabrication plant. The HTR-PM will have two 105 MWe modules connected to a single steam turbine. Plant design commenced in 2001. The project was initially supposed to be finished by around 2013 (Zhang et al. 2009). However, construction did not begin until the end of 2012, at which point the plant was projected to begin supplying electricity to the grid by 2017 (Dalton 2013). But fuel fabrication did not commence until 2016, and the project has been beset by further delays. Apparently China's initial plan was to connect the first module to the grid before completing the second module (Jian et al. 2014), but at some point it decided to defer commissioning until the second module was also built. It is unclear why the plan was changed. The plant is currently projected to be fully operational by the end of 2022.

Despite the delays in the HTR-PM project, China continues to plan construction of an even larger, 600 MWe version, the HTR-PM600, which would have six modules connected to a single steam turbine.

In the United States, a number of startup companies are again attempting to commercialize HTGRs. X-Energy has revived the pebble-bed small modular reactor design that was abandoned by Eskom in 2010. Its Xe-100 reactor would generate 200 MWth (75-80 MWe) and would be bundled in 300 MWe "four-packs" connected to a single steam turbine (X-Energy n.d.).¹⁹ In October 2020 the DOE chose the Xe-100 four-pack for commercial demonstration by 2027 under the ARDP, after previously awarding X-Energy a number of grants, including \$8.9 million to support design and Nuclear Regulatory Commission (NRC) licensing of a TRISO fuel fabrication facility.

Framatome (formerly New NP, a subsidiary of Areva NP) designed a prismatic-block HTGR that was chosen by the NGNP industry alliance in 2012 as the optimum NGNP design before the DOE terminated the project. The NGNP Alliance was unable to find another customer for the reactor. More recently, Framatome submitted information about its HTGR design for evaluation in the DOE's demonstration and test reactor study (Petti, et al. 2017). However, there is no indication that Framatome is pursuing further development of its HTGR at this time.

General Atomics, which heavily promoted a modular thermal HTGR design in the past, now appears focused on gas-cooled *fast* reactors. In 2009, it began to develop the Energy Multiplier Module (EM2), a 265 MWe convert-and-burn design that in theory would operate for up to thirty years on one load of fuel. The project has not progressed beyond the research phase. More recently, General Atomics and Framatome announced the start of a collaboration to develop a 50 MWe gas-cooled fast reactor that would operate on a more practical nine-year refueling cycle (GA 2020). However, the DOE considers gas-cooled fast reactors, which have never been demonstrated, to be the least technologically mature of all the NLWRs it evaluated (Petti et al. 2017), and the present report does not discuss them further.²⁰

More recently, TRISO-fueled HTGRs have gotten an additional boost from the Department of Defense (DOD), which issued a "request for solutions" in 2019 for the "first phase of a multi-phase prototype project for a small mobile nuclear reactor," with a capacity of 1 to 10 MWe—otherwise known as a microreactor. Decades after ending its previous nuclear power program, the Army is now interested in developing microreactors for domestic bases in remote locations and for forward operating bases overseas (Lyman 2019). The request for solutions specified that these reactors must use TRISO fuel. In March 2020, the DOD awarded contracts to three teams—BWXT, Westinghouse, and X-Energy—to begin design work on a prototype (DOD 2020). The DOD will decide whether to proceed with actual prototype construction after a two-year "design maturation period." X-Energy, which received the largest award (\$14.3 million), will have to scale its Xe-100 down by a factor of around 10 in order to provide a prototype microreactor meeting DOD requirements.

Two foreign companies, StarCore and U-Battery, are pursuing small modular HTGRs and microreactors. StarCore would use a static pebble-bed core, and U-Battery a prismatic core. StarCore submitted a microreactor proposal to DOD but was not selected. U-Battery, a 10 MWth microreactor, which is a project of a consortium led by the URENCO uranium enrichment conglomerate, hopes to have a demonstration plant operating by 2028 (U-Battery 2019).

Safety

HTGRs have several attractive characteristics from a safety perspective. First, similar to LWRs, they have negative temperature coefficients, so that reactor power will tend to decrease if the fuel overheats (see chapter 2). This is an important inherent safety feature. Second, the core has a low power density because of the presence of a large mass of graphite, so that it tends to heat up slowly if cooling is interrupted. The third safety feature is the ability of the TRISO fuel itself to retain fission products during both normal and certain accident conditions.

It is this latter characteristic that HTGR developers most often emphasize when claiming that their reactors are safer than LWRs. For example, the NGNP Industry Alliance writes that

the high temperature and robust structural capabilities eliminate concerns of fuel damage that could lead to significant release of radioactive materials from the nuclear fuel. The ceramic-coated nuclear fuel provides the primary containment for radioactive materials rather than depending on a containment building (NGNP 2010).

Or, in the words of a leading HTGR expert, David Petti, the fuel is the “sine qua non” of the design (ACRS 2011). The Pentagon was apparently swayed by these arguments in specifying that the microreactors it is considering for deployment should only use TRISO fuel, citing its “robust safety” and even “minimized fission product release due to blast conditions/adversary attacks” (DOD 2019).

Based on the claim that each TRISO fuel particle has its own “containment,” current HTGR designs do not enclose the reactor vessel in a leak-tight, high-strength containment building, which is standard for LWRs. Instead, the designs call for a less protective “confinement” building with filtered exhaust systems. (This aspect is no doubt a desirable feature for DOD’s mobile microreactor program.) The safety of HTGR designs without a conventional containment structure depends critically on the fuel performing as advertised to contain fission products.

Thus, there are two overarching questions relevant to HTGR safety. The first is whether properly manufactured TRISO fuel is actually capable of retaining fission products to the extent necessary to adequately protect public health and the environment without the need for a leak-tight containment, during both normal operation and accidents or sabotage. The second is the extent to which TRISO fuel can be reliably manufactured according to its design specifications. No matter how safe the fuel design, a high defect rate would undermine its performance.

While these fuel concerns generally apply to all HTGRs, the safety challenges of pebble-bed reactors with constantly moving fuel are more serious. The complexity inherent in such a system makes it more difficult to monitor the condition of the reactor core and to predict the performance of the fuel pebbles. (These problems are even more pronounced in reactors with fluid fuel, such as MSR, which are discussed in the next chapter.)

Finally, as discussed below, other HTGR design features—such as the use of graphite, which chemically reacts with air and water (oxidizes)—introduce safety problems that are not issues for LWRs.

TRISO FUEL PERFORMANCE UNDER ACCIDENT CONDITIONS

The primary barrier to a release of radioactivity from an HTGR during an accident is the TRISO fuel itself. However, as is the case for any nuclear fuel, TRISO fuel particles will lose their integrity when heated above a certain temperature and will release fission products. Although the particles are embedded in fuel elements, some of these fission products can be transported through the fuel element graphite matrix and end up in the helium coolant. Depending on the nature of the accident, fission products could then be released into the confinement building and eventually into the environment. Therefore, a detailed understanding of the maximum temperatures that can occur during HTGR accidents and the behavior of TRISO fuel under those conditions is critical for assessing HTGR safety.

As noted above, under normal operating conditions, the maximum coolant temperature in an HTGR is about 800°C. However, if normal cooling is disrupted, this temperature could greatly increase. The highest temperature at which properly manufactured TRISO fuel has been observed to fully maintain its integrity is around 1600°C. Above this temperature, the picture is a lot less clear. At higher temperatures, TRISO fuel particles have been observed to release substantial quantities of fission products, but experimental data in this area are limited and fundamental mechanisms are not well understood (Demkowicz, Petti, and Gougar 2017).

Therefore, to demonstrate HTGR safety, it is necessary to show that (1) the peak fuel temperature will not exceed 1600°C for design-basis accidents, and (2) if the fuel does exceed 1600°C during a beyond-design-basis accident, fission product releases to the environment will not significantly impact public health and safety. (Decisions on where to draw the line between design-basis and beyond-design-basis accidents typically would be up to the regulator.)

A number of different accidents, such as a reactivity excursion, could cause the fuel to overheat and potentially exceed 1600°C. The most challenging design-basis accident

is usually assumed to be a rapid loss of helium coolant. If a large breach occurs in the primary coolant system, then the pressurized helium gas coolant would quickly escape and the system would depressurize. Unlike LWRs, HTGRs do not have systems to pump emergency supplies of coolant into the reactor vessel if a loss of coolant accident occurs. And, as will be explained below, water could not be injected either.

If such a depressurization accident occurs, there are two alternative safety approaches to keep the fuel temperature below 1600°C. One is to use an external active (e.g., motor-driven) coolant system to remove decay heat from the surface of the reactor vessel through forced convection. The other is to rely only on passive means: natural convection cooling. For the latter approach to work, the maximum power rating and physical size of the reactor must be limited.

Partly because of the problems that larger HTGRs such as Fort St. Vrain experienced with active decay heat removal systems using forced convection, current HTGR designs are limited to below about 300 MWe so that they can rely on passive means alone for emergency cooling. These designs, such as the Xe-100, are being marketed as small modular reactors.²¹ However, even the reference 200 MWth (75 MWe) design for the Xe-100 could reach a temperature of over 1700°C during a depressurized loss of coolant accident and thus fails to meet the 1600°C peak temperature limit (Mulder and Boyes 2020). To meet this fundamental safety criterion, the power rating would have to be reduced to 165 MWth, a smaller and even less economical reactor (Mulder and Boyes 2020).

UNCERTAINTIES IN ALLOWABLE PEAK TEMPERATURES

Even with a limitation on the reactor power and passive cooling, it turns out to be difficult to prove that the 1600°C temperature limit will not be exceeded for a depressurization accident. This is because HTGR accident analyses have large uncertainties, and peak temperature calculations are imprecise. It is difficult to validate the computer models used for these analyses because key parameters, such as core temperatures, cannot be directly measured.

HTGR operational experience has shown that hot spots can develop that greatly exceed the maximum temperatures predicted by models (Carlson 2014). One reason is the occurrence of so-called bypass flows, which are unpredictable changes in the flow of the gaseous coolant due to random structural changes in the core (Beck and Pincock 2011). Although these challenges exist for both prismatic-block and pebble-bed reactors, modeling is particularly problematic for moving-fuel pebble bed reactors, which have hundreds of thousands of circulating pebbles.

These uncertainties were observed at the AVR (Arbeitsgemeinschaft Versuchreaktor) pebble-bed test reactor

in Germany. Researchers there calculated a maximum fuel temperature of 1070°C. They devised a method to determine whether the temperatures of individual pebbles would exceed safety limits by inserting wires with a range of melting points up to 1280°C into unfueled monitor pebbles (Moorman 2009). Depending on the number of wires that melted, operators could determine the maximum temperature—as long as the peak melting temperature was not exceeded. But in a significant number of monitor pebbles all of the wires had melted, indicating the coolant temperature had exceeded 1280°C—more than 200 degrees above 1070°C (Moormann 2009). It was estimated afterward that the actual peak core temperature could have reached 1420°C, or 350 degrees above the calculated maximum.

This inability to predict peak fuel temperatures is a problem because significant quantities of fission products could be released if actual fuel temperatures exceed the 1600°C safety limit. An accident analysis that calculates a peak fuel temperature of 1600°C could underestimate the true peak by hundreds of degrees.

UNCERTAINTIES IN FISSION PRODUCT RELEASES AT HIGH TEMPERATURES

Moreover, there are limited data about the performance of TRISO fuel at temperatures of 1600°C and above. Recent testing conducted as part of the DOE's TRISO fuel development program has demonstrated that the fission product releases from TRISO fuel as a function of temperature are complex and still not fully understood (Hunn et al. 2017a, EPRI 2020). In these tests, TRISO fuel compacts are irradiated in the Advanced Test Reactor at Idaho National Laboratory and then heated outside of the reactor to simulate accident conditions. The tests found that some irradiated TRISO fuel compacts had relatively high releases of certain fission products even at 1600°C. For example, fractional releases of strontium-90 and europium-154, two long-lived and radiotoxic fission products, were close to 10 percent at 1600°C for one uranium oxycarbide TRISO fuel compact. These release fractions are comparable to or larger than typical releases from light-water reactor fuel in a design-basis accident.

For some uranium oxide compacts, cesium releases reached nearly 1 percent at 1600°C and 10 percent at 1700°C before the experiment was prematurely terminated—due to the unexpectedly high release (EPRI 2020). This is comparable to an LWR release to the environment in a severe accident. Releases of cesium from the Fukushima accident through breached containment buildings are estimated at a few percent of the core inventory. Therefore, these new data do not support the claim that fission product releases from TRISO fuel will always be so low that a containment building is not necessary.

One troubling gap in TRISO fuel accident performance data is the lack of experimental observations of the behavior of radioactive iodine. Radioactive iodine fission products are one of the most significant contributors to off-site radiation doses in a severe accident. However, because these isotopes are predominantly short-lived (iodine-131, with a half-life of eight days, is one of the longest-lived), by the time irradiated TRISO fuel is ready for accident testing, the iodine has largely decayed away. Until sufficient data on iodine release have been collected, it will be premature to make conclusions about whether HTGRs need robust containments and off-site emergency planning measures.

TRISO FUEL FABRICATION AND QUALIFICATION

As noted above, the safety of HTGRs depends heavily on the quality of the fuel. TRISO fuel performance is very sensitive to slight imperfections that could occur during the complex manufacturing process. To meet HTGR safety goals, the production of TRISO fuel particles requires very high standards of quality control for multiple design parameters to maintain an acceptably low defect rate—below 1 in 100,000 for some parameters (Petti et al. 2010). In the words of DOE experts:

The required level of fuel performance and fission product retention reduces the radioactive source term by many orders of magnitude relative to source terms for other reactor types and allows a graded approach to emergency planning and the potential elimination of the need for evacuation and sheltering. Achieving this level, however, is predicated on *exceptionally high* coated-particle fuel fabrication quality and *excellent performance* under normal operation and accident conditions [emphasis added] (Petti, Collin, and Marshall 2017).

However, the United States has not yet demonstrated that it can produce fuel of “exceptionally high . . . quality” that exhibits “excellent performance.” The historical performance of TRISO fuels in US test and demonstration reactors was far less successful than the experience in Germany. In particular, US-fabricated fuel released fission product gases at a rate 1000 times greater than German-fabricated fuel during normal operation (Petti et al. 2002). Although the irradiation conditions were different in the United States and Germany, the disparity is believed to be primarily due to the lower defect rate of the German TRISO fuel particles, which was on the order of 100 out of 3.3 million particles fabricated (a rate on the order of 3.3 per 100,000, which still does not meet the current safety specifications).

The DOE initiated a comprehensive Advanced Gas Reactor (AGR) Fuel Development and Qualification Program in

2002 to address the problem of poor-quality US fuel production, taking advantage of the German experience. The goal was initially to develop fuel for the Generation IV Very High Temperature Reactor and NGNP projects, but after those ended, its focus shifted to qualify TRISO fuel and establish a US commercial fuel vendor. This program, originally scheduled to be completed in the mid-2020s, was to culminate in a series of formal fuel qualification irradiation tests to provide sufficient data “to demonstrate compliance with statistical performance requirements (AGR-5,6) as well as a test at elevated temperature to establish safety margin (AGR-7)” (Marshall 2019). An eighth test to validate fission product transport models, AGR-8, was cancelled. All irradiation tests were conducted in the Advanced Test Reactor at Idaho National Laboratory.

However, these tests have encountered a number of technical problems, including many failures of thermocouples—instruments needed to accurately measure the very high temperatures at which the TRISO irradiations were conducted. The AGR-2 test, which began in 2010, lost most of its thermocouples early on and experienced other problems that rendered critical fuel performance data useless after the third irradiation cycle, or only about one-quarter of the way through the test (EPRI 2020). And in late 2018, a high rate of thermocouple failures and additional technical problems plagued AGR-5/6/7, the final—and most important—test series, including cracks in an irradiation capsule and plugs forming in the outlet gas lines (Palmer 2019). Three of the irradiation capsules containing fuel for the formal qualification tests eventually lost all functioning thermocouples, so that critical temperature data were not obtained for a number of irradiation cycles. And in 2019, one of the capsules suddenly began releasing fission products at a high enough rate to exceed the yearly dose limit for operation of the Advanced Test Reactor, requiring operators to isolate the capsule and stop collecting fission product release data from it (Pham et al. 2020). There is little public information about what impact these problems will have on completion of the program.

Moreover, the United States has still not achieved production of TRISO fuel that meets all specifications. For example, a batch of TRISO fuel supplied by BWXT to the DOE in 2016 for potential use in the AGR program was rejected because it failed to meet the 1-in-10000 specification for defects in one of the coated layers (Hunn et al. 2017b). Subsequent TRISO fuel lots produced by BWXT at a “near-commercial scale” for the AGR-5/6/7 tests did not meet multiple specifications, but were deemed acceptable for the test. Consequently, the fuel has performed worse than the first test, called AGR-1, which used fuel that was produced only at a laboratory scale (Pham and Scates 2019).

Therefore, many years of additional process development and irradiation testing in the United States will likely be necessary to ensure that fuels can be consistently manufactured to the very high standard required and to meet performance goals.

Few data are available on the fuel performance of the only operating HTGR today, the 10 MWth HTR-10 pebble-bed test reactor in China. TRISO fuel for the HTR-10 was produced in a pilot fabrication line at Tsinghua University in the 1990s. It has been reported that some fuel batches experienced performance problems during testing in a Russian test reactor that could be traced to manufacturing defects (Beck and Pincock 2011). However, the tests were apparently not well instrumented. In one heating test, six percent of the particles failed, but researchers do not know what the maximum temperature actually was, although they believe that it exceeded 1600°C (Tang et al. 2006).

Subsequently, China built a larger pilot production plant at Tsinghua to fabricate sample fuel for the HTR-PM reactor. More recent data on this fuel, which was irradiated in the Petten reactor in the Netherlands and sent to the Joint Research Center in Karlsruhe, Germany, for accident testing, confirm that fission product releases are fairly low at around 1600°C but greatly increase as the fuel temperature is increased, with cesium-137 fractional release of more than 5 percent at 1770°C (Freis et al. 2020).

While these fuel qualification tests were being conducted in Europe, China built a larger plant for HTR-PM fuel fabrication at the reactor site in the city of Baotou and began operating it in 2016, before the tests were even completed.

OTHER HTGR HAZARDS

Fuel heat-up during a depressurization accident is not the only mechanism that could cause fuel damage and fission product release. If air or water leaks into the reactor core (referred to as “air or water ingress”), the consequences could be severe. An HTGR contains a large quantity of graphite both in core structures and in the fuel itself. Graphite can undergo energy-releasing oxidation reactions if exposed to air or water, causing it to lose mass and weaken. It can also react with water to form flammable gases. In addition, water ingress can cause other serious problems, including a rapid increase of reactivity (given that water is a good moderator of neutrons).

Therefore, both air and water must be prevented from entering the core to a very high degree. If depressurization occurs, a significant quantity of air could enter the primary coolant circuit, coming into contact with and oxidizing the graphite fuel elements and structural materials. Operator

errors could enable such an accident (Carlson 2014). Water and/or air ingress is suspected as the cause of the large number of TRISO particle failures and subsequent high fission product releases that occurred in one of the capsules during the AGR-5/6/7 irradiation at the Advanced Test Reactor at the INL—high enough to exceed a downwind annual radiation dose limit (Palmer 2020; Pham et al. 2020).

While there is no question that graphite will undergo combustion when exposed to air, especially at high temperature, there is a long-running debate about whether high-purity, reactor-grade graphite can ever undergo self-sustaining combustion—that is, to actually “catch fire.” Self-sustaining combustion means that the heat of the reaction itself is sufficient to maintain the process. Although most assessments of the 1986 Chernobyl accident describe the burning of the graphite moderator as a “graphite fire,” some HTGR researchers dispute this terminology and go as far as to assert that “self-sustained oxidation is physically impossible in nuclear grade graphite” (Windes et al. 2014). However, other analysts are not willing to make such unequivocal conclusions, conceding that self-sustaining oxidation reactions during air ingress can occur “in extreme situations” (Morris et al. 2004) or are merely “difficult to achieve” (Areva 2010). The French nuclear safety research organization Institut de Radioprotection et de Sûreté Nucléaire (IRSN) concluded that “the risk of a graphite fire in such [water or air ingress] conditions cannot be ruled out,” and that “even if the [graphite] degradation is only localized, it may still have significant consequences” (IRSN 2015). Thus, whether graphite oxidation at high-temperature during an HTGR accident is self-sustaining or driven by high fuel temperature, it is a safety issue that must be rigorously assessed.

The potential for water to enter the core if there were a severe flood at the reactor site, such as the tsunami that triggered the Fukushima meltdowns, clearly needs to be addressed. Moreover, even if the risk of accidents that could result in air or water ingress is shown to be small, the potential for sabotage will always be present. For this reason, HTGRs will require robust security.

For pebble-bed reactors, another source of radioactive material is graphite dust produced by friction between the pebbles, a phenomenon that has not been accurately modeled (Humrickhouse 2011). (Prismatic-block HTGRs are believed to generate far less dust, according to the IRSN in France.) The graphite becomes radioactive both from absorbing fission products that are released from the TRISO fuel during normal operation and through neutron irradiation of its constituent elements (for instance, non-radioactive carbon can absorb neutrons and become radioactive carbon-14). This dust can be expelled in the event of a primary coolant depressurization event, resulting in a significant release of radioactivity

from the core even if the reactor fuel remains intact.

Another issue is that graphite swells when irradiated, which can cause a variety of problems, especially since it provides support structures for HTGR cores.

CONTAINMENT AND EMERGENCY PLANNING REQUIREMENTS

As discussed in chapter 2, nuclear plant safety is rooted in the concept of layers of protection known as “defense-in-depth.” For current-generation LWRs, the NRC’s defense-in-depth requirements include a robust, leak-tight containment building to limit the release of radioactive material to the environment in the event of a core melt. In addition, the NRC requires current LWRs to have emergency preparedness programs, including the designation of emergency planning zones with a 10-mile radius around each plant site. These are areas in which measures such as evacuation and potassium iodide distribution could be carried out if an accident or terrorist attack caused an off-site radiological release.

However, HTGR developers and the DOE have long argued that these two defense-in-depth measures—leak-tight containments and emergency planning zones—are not necessary because each TRISO fuel particle has its own containment and thus there is virtually no risk of a release of a significant quantity of fission products from the fuel (Lyman 2001). Such arguments were viewed skeptically by the NRC for many years, but recently they have gotten traction at the agency, not only for HTGRs but for all NLWRs and even for small modular LWRs.

In 2018, the NRC commissioners unanimously approved a staff proposal to develop “functional containment” performance criteria that would allow relaxation of the current requirement for a leak tight, pressure-resisting containment structure by taking credit for other design features such as the use of TRISO fuel (Vietti-Cook 2018). That would pave the way for NRC approval of HTGRs with filtered, vented confinement buildings. Also, in December 2019, in a 3-1 vote, the commissioners approved publication of a draft rule that would allow NLWR and small modular LWR applicants to reduce or eliminate emergency planning zones based on off-site dose calculations crediting features such as TRISO fuel (NRC 2019).

Rainer Moormann, a German HTGR researcher who has become a leading skeptic of the technology, concludes that future pebble-bed HTGRs should include leak-tight containments, given the many unresolved safety issues including the potential for fuel temperatures and fission product releases to greatly exceed expected values (Moormann 2009). In a recent critique of the HTR-PM commercial demonstration pebble-

bed reactor that is under construction in China, Moormann and collaborators proposed a number of safety upgrades to compensate for the absence of a leak-tight containment at the reactor, such as improving the confinement vent filtration system (Moormann, Kemp, and Li 2018). However, even such upgrades cannot provide the same level of safety assurance as a robust containment.

In addition, off-site emergency planning is critical even for reactors with conventional containments, because they can fail—as the 2011 Fukushima accident demonstrated. Removing one layer of defense-in-depth for a reactor with unproven safety features is risky enough—removing two layers is even more reckless.

Thus, given the uncertainties in the performance of TRISO fuel and other HTGR safety issues, such as graphite dust generation, there is insufficient justification for eliminating robust containments and off-site emergency planning zones for HTGRs. Much more work will be required to achieve the necessary level of assurance. Unless the HTGR’s safety basis can be fully validated through testing that covers the full range of severe accident and terrorist attack scenarios, it would be unwise for the NRC to license HTGRs without all the layers of protection that reactors now rely upon to protect the public.

Sustainability

Sustainability is one area where HTGRs have clear disadvantages compared to LWRs. HTGRs use uranium less efficiently and generate a greater volume of nuclear waste.

URANIUM UTILIZATION EFFICIENCY

The more robust TRISO fuel can likely achieve substantially higher burnups than LWR fuel, which could in principle lead to better uranium utilization (see chapter 3), although this will require further fuel qualification activities to establish. However, this fuel requires a higher level of uranium enrichment in order to do so—from seven to 19.9 percent uranium-235 (U-235), depending on the reactor design. The net result is that more natural uranium must be enriched to generate a given amount of power, reducing the uranium utilization efficiency (Bays and Piet 2010).

As noted above, the higher coolant temperatures of HTGRs result in a thermodynamic efficiency up to one-third higher than that of LWRs. Even so, this is not enough to overcome the penalty resulting from the additional uranium needed to produce the higher-assay LEU fuel. A prismatic HTGR with 38 percent thermal efficiency and a fuel burnup twice that of an LWR would be about 70 percent as uranium-efficient

(Piet, Bays, and Soelberg 2010).

Pebble-bed HTGRs with mobile fuel are about one-third more uranium-efficient than static-fuel designs. The ability to continuously refuel the core allows for more flexible fuel management. Nevertheless, they are still slightly less uranium-efficient than LWRs (Piet, Bays, and Soelberg 2010). The X-Energy Xe-100, even with an assumed average fuel burnup of 164,000 MWd/MTHM, or 16.9 percent—about three times that of a typical LWR—would still require slightly more natural uranium as an LWR per gigawatt-year. However, it is unlikely that X-Energy’s TRISO fuel will be able to achieve such a high burnup, at least in the near term. Because the AGR-5/6/7 tests previously described at Idaho National Laboratory were not completed (Pham et al. 2020), the peak fuel burnup achieved was likely only about 15.3 percent, or 148,000 MWd/MTHM, falling short of the Xe-100 target burnup and the original test goal of 19.1 percent (Palmer 2019). The NRC is not likely to approve a peak burnup greater than—or even close to—this value for the planned Xe-100 demonstration plant. Consequently, the uranium usage will very likely be greater than initial expectations.

Researchers have studied ways to make pebble-bed HTGRs more uranium-efficient, without success. A study using a sophisticated computer algorithm to optimize the design found that increasing the complexity of pebble-bed fueling strategies did not significantly reduce uranium demand (Tavron and Shwageraus 2016). And the use of thorium in addition to low-enriched uranium (LEU), which could in principle reduce the need for uranium ore, does not seem to help either. Studies have found that adding thorium did not significantly increase the efficiency and could even decrease it (Xia et al. 2014; Tavron and Shwageraus 2016).

Increasing the thermal efficiency is another way to improve the uranium utilization of HTGRs. Replacing the conventional steam cycle with a helium-based Brayton cycle—which is still at a low level of technical maturity—could in theory increase efficiency to near 50 percent. Even so, improvements would be modest.

WASTE

HTGRs do not offer significant advantages with respect to nuclear waste generation compared to LWRs. HTGR fuel consists of TRISO fuel particles embedded in a carbonaceous matrix, and because the uranium is diluted in a large mass of non-fuel material, HTGR spent fuel has a low power density compared to LWR spent fuel. However, this dilution results in a 10-fold increase in waste volume per unit of electricity generated compared to LWRs (Lyman 2001).

As discussed in chapter 3, the high-level waste volume

reduction that is achieved by spent fuel reprocessing does not generally increase the capacity of a geologic repository because decay heat load is typically the limiting factor, not waste volume. In other words, it may not be possible to cram packages of more concentrated high-level waste closer together if the waste is hotter than spent fuel. However, this argument does not necessarily apply in the opposite direction. Depending on the detailed characteristics of a repository, it may not be possible to dispose of a more dilute waste form than LWR spent fuel in the same amount of space because of physical limitations. Also, the increased potential for a criticality accident—an inadvertent chain reaction—in a repository given the greater uranium enrichment of HTGR fuels would have to be taken into account.

Partly because HTGRs use HALEU fuel with a lower concentration of U-238 than LWRs, they generate approximately one-half as much plutonium and other TRU per GW-year (Shropshire and Herring 2004). However, this is not a significant reduction.

The impact on repository capacity is not the only consideration with management of HTGR spent fuel—there are also challenges associated with storage and transport of such a large volume of waste. More waste packages would be required to dispose of a given amount of uranium, requiring more materials, more shipments, and increasing cost. A 2015 European Commission report concluded that “the direct disposal of spent [HTGR] fuel would possibly not be acceptable in case of a larger . . . fleet, because of the large associated volumes and large amounts of steel for the containers” (Knol et al. 2015).

Also, the large amount of irradiated carbon in the waste contains a significant inventory of the long-lived radioactive isotope carbon-14, which would contribute substantially to the repository’s radioactive release to the environment. This would be particularly troublesome in a repository above the water table such as Yucca Mountain, because the carbon-14 could be released in the form of carbon dioxide to the atmosphere, allowing it to spread widely.

Due to these and related issues, it is far from clear whether it would be safe or practical to directly dispose of HTGR spent fuel in a geologic repository. If not, then a method would have to be devised to separate the carbonaceous fuel matrix from the fuel particles before disposal and reduce its volume (Li, Ma, and Wang 2014). The feasibility and cost of such methods have not been determined. Whether done mechanically or chemically, however, the residual carbonaceous material would also be radioactive waste, although it would likely be classified as low-level waste that could be disposed of in a less robust facility than the irradiated TRISO fuel would require.

Proliferation/Terrorism Risk

The proliferation risks of HTGRs and their associated fuel cycles would depend on the type of the reactor, the characteristics of the fuel, and whether the spent fuel would be stored for eventual direct geologic disposal or reprocessed.

As discussed in chapter 4, for any reactor type, fuel cycles involving reprocessing pose greater proliferation risks than once-through cycles because of the risks of diversion of weapon-usable materials such as separated plutonium. Multiple diversion scenarios must be considered, including diversion from reprocessing and plutonium fuel fabrication plants. However, even if a country has no declared reprocessing plants, there is still a risk that spent fuel could be diverted from a reactor to a covert reprocessing plant. Therefore, International Atomic Energy Agency (IAEA) safeguards must always be applied to reactors to verify that spent fuel is not diverted, although the inspection goals may be less stringent in countries with no declared reprocessing plants.

Thus, in assessing the proliferation risks of HTGRs compared to LWRs on a once-through cycle, a major factor is how effectively safeguards can be applied to the reactors themselves.

Another factor is the proliferation risk posed by the fresh fuel. As discussed above, to take advantage of the high burnup potential of TRISO fuel, HTGRs must utilize high-assay LEU (HALEU). For instance, the Framatome prismatic-block and X-Energy pebble-bed HTGRs would both use many tons of 15.5 percent-enriched LEU. As discussed in chapter 4, although this material is impractical for direct use in a weapon, it has a Category II security ranking and is more attractive for illicit use than are lower-enriched LEU fuels. The use of HALEU is particularly problematic in fabricating TRISO fuel on a commercial scale, because of the increased difficulty in accounting for the huge number of TRISO fuel particles that would be produced annually. A pebble-bed HTGR would require a supply of around 10 billion TRISO particles per gigawatt-electric-year (GWe-year), compared to a few million uranium fuel pellets per GWe-year for an LWR.

Another consideration is whether the special physical and chemical properties of HTGR fuel would make it less attractive for diversion for weapons use than LWR fuel. Some argue that HTGRs are more proliferation-resistant than LWRs because it is more difficult to reprocess their spent fuel. In order to do so, spent TRISO fuel particles would have to be separated from their carbonaceous matrices, and the robust particle coatings would have to be breached in order to extract fissile materials from the fuel kernels. While such techniques were demonstrated on a pilot scale in the 1980s, there has been no demand to develop the reprocessing

processes on an industrial scale, since HTGRs have not been commercially deployed.

After reviewing this issue in 2006, DOE researchers found that although HTGR spent fuel would be technically challenging to reprocess, it would still be attractive to potential proliferators as a source of material for nuclear weapons (Durst et al. 2009). A 2010 Areva (now Framatome) study similarly concluded that:

[t]hough the presence of strong coatings on fuel particles adds a difficulty for retrieving the fissile content of the fuel, it is not impossible: this is done currently by mechanical processes for recovering fabrication scraps, admittedly not with irradiated fuel, and new processes in development ... might make it even easier (Areva 2010).

The absence of an industrial-scale HTGR spent fuel reprocessing infrastructure does not mitigate the proliferation risk, because even a small-scale clandestine facility might be able to separate one significant quantity of plutonium within a year. Therefore, one cannot conclude that HTGRs would require less stringent safeguards than LWRs by virtue of their fuel. Moreover, given the potential issues with direct disposal of HTGR spent fuel discussed above, it is likely that there would be renewed interest in developing industrial-scale reprocessing if there were deployment of a large HTGR fleet.

HTGR SAFEGUARDS

A key factor in assessing the proliferation risk posed by an HTGR is the way in which it is refueled. Prismatic-block HTGRs would be refueled in a similar manner to LWRs. The reactor would have to be shut down, the fuel blocks loaded and unloaded in batches, and the reactor vessel resealed for the next operating cycle. Keeping track of HTGR prismatic fuel blocks, which are readily countable items, would be no more challenging than keeping track of LWR fuel. Because HTGR fuel can achieve higher burnups than LWR fuels, the operating cycle would be longer and refueling less frequent, which could make prismatic-block HTGRs somewhat easier to safeguard than LWRs. However, since current HTGR designs are small modular reactors, a commercial HTGR plant with the same generating capacity as a single large LWR would have multiple modules that would be refueled at different times. This would likely require more visits from IAEA inspectors, increasing the cost of safeguards.

In contrast, pebble-bed HTGRs have characteristics that would make them more difficult to safeguard than LWRs. First, the reactor would be continuously refueled while operating, providing greater opportunities for diversion of both fresh and irradiated fuel. While this problem is already encountered at on-line-fueled reactors such as the Canadian-

designed CANDUs, a pebble-bed reactor would pose greater challenges. Instead of the thousands of fuel bundles in an CANDU core, the core of a single pebble-bed reactor module would contain hundreds of thousands of fuel pebbles. The presence of multiple small reactor modules instead of a single large reactor would increase the number of items on site and the complexity of applying safeguards at such a facility.

For example, the core of each Xe-100 76 MWe pebble-bed reactor module would contain 220,000 pebbles. Assuming that the fuel can actually achieve the design average burnup of 163,000 MWd/MTHM, about 1060 pebbles would be loaded into and discharged from the reactor each day—a flow rate of one every 80 seconds. Of these, 175 fresh fuel pebbles would be added and 175 spent fuel pebbles would be discarded to waste storage daily. Spent fuel storage bins would hold hundreds of thousands of pebbles each. It would take a significant effort to accurately keep track of this huge number of fuel pebbles and the nuclear material they contain. The small size of the pebbles would also make them easier to conceal and steal (IAEA 2014). Also, even for very high burnups, the spent pebbles do not have the “self-protecting” radiation field characteristic of LWR spent fuel, because each one has only a very small quantity of fission products (Chung et al. 2012).

On the other hand, the amount of enriched uranium fuel in each fresh fuel pebble and the amount of plutonium in each spent fuel pebble is low—around 7 grams and 0.12 grams, respectively, for the Xe-100. Tens of thousands of pebbles, or around 20 percent of a single Xe-100 core—would be needed to acquire enough material for a nuclear weapon. An abrupt diversion of this much fuel would likely be observable. Nevertheless, safeguards inspectors would need to be able to detect multiple small diversions that could result in the accumulation of a significant quantity of fissile material over time—again, complicated by the number of operating units and the spent fuel storage bins at a site.

At LWRs or CANDUS, every fuel assembly can be uniquely identified by an engraved serial number, which enables it to be tracked throughout the facility and be verified by IAEA inspectors. While that is not an option for graphite-based fuel pebbles, a technical alternative has been proposed for an internal identifier (Gitau 2011), although the proposal has not gone beyond the conceptual stage. However, even if it were possible to uniquely identify each pebble, it would not be practical to use the identifier to track each item throughout the facility in real time. DOE researchers have pointed out that “the declaration and accounting of such large numbers of fuel pebbles individually would be onerous” and concluded that existing reactor safeguards approaches would not work for pebble-bed reactors (Durst et al. 2009).

Instead, the DOE researchers proposed a new approach that more closely resembles safeguards at bulk-handling facilities (such as nuclear fuel fabrication plants) than those applied at traditional reactors. However, as discussed in chapter 4, the accounting procedures at such facilities have inherent uncertainties, making it harder to distinguish genuine diversions from statistical and measurement errors. Consequently, a non-zero “material unaccounted for” is to be expected at such facilities. Such uncertainties have already been reported at the HTR-10 in China, where there was “uncertainty about the precise number of pebbles in the core, because the redundant facility pebble-counters did not exactly agree” (Durst et al. 2009). This problem would be compounded in a commercial-scale reactor such as the Xe-100, which would have 10 times as many pebbles as the HTR-10.

Therefore, instead of trying to count and keep track of every individual fuel pebble, operators would use radiation-based fuel flow monitors, designed to detect anomalies in the pebble streams (Durst et al. 2009). Such systems would have to be sensitive enough to distinguish signals caused by fuel diversions from normal statistical variations, and would also have uncertainties.

As a result, inspectors would have to supplement these techniques with containment and surveillance measures, which are inherently less reliable than material accounting. If such measures were to be lost temporarily—for instance, if a surveillance camera stopped working—then the only way to recover continuity of knowledge would be to conduct a time-consuming inventory of all material at the facility. Also, in contrast with LWR spent fuel, safeguards inspectors would not be able to directly observe pebble-bed reactor spent fuel in storage (Durst et al. 2009). At LWRs, spent fuel is stored below several meters of water, providing radiation shielding but allowing inspectors to view the spent fuel. HTGR spent fuel pebbles cannot be stored in water but must be immediately sent to dry storage bins, where they cannot be viewed.

STATUS OF PEBBLE-BED SAFEGUARDS APPROACHES

In order to implement the DOE researchers’ proposed safeguards approach, the IAEA would have to develop new criteria and technologies, which could take many years and hefty resources (Durst et al. 2009). At the time of the proposal, Eskom was actively pursuing deployment of the pebble-bed modular reactor in South Africa, and other non-nuclear weapon states had expressed interest in the technology. However, Eskom suspended work on this reactor in 2010, and the IAEA decided not to pursue development of pebble-bed reactor safeguards approaches, given the lack of interest.

But current pebble-bed developers are again hoping to export their products around the world to non-nuclear

weapon states where IAEA safeguards would be required. China National Nuclear Corporation signed a memorandum of understanding with Saudi Arabia to conduct a reactor feasibility study in 2017 and has had discussions with Indonesia and Egypt (Nicobar Group 2017), although no projects appear to have materialized. More recently, X-Energy signed a letter of intent with Jordan in November 2019 to build four Xe-100 units by 2030. If these projects move forward, the IAEA will soon need a workable and effective safeguards approach. The IAEA reported in 2019 that it was working with China to develop safeguards approaches for the HTR-PM under China's voluntary offer agreement, but no details are available (IAEA 2018). The United States could assist the IAEA in developing pebble-bed safeguards approaches by offering the proposed Xe-100 commercial demonstration plant for IAEA safeguards during the design phase—and providing funding for the effort.

Readiness for Commercial Demonstration and Near-Term Deployment

As discussed in chapter 1, the 2017 DOE advanced demonstration and test reactor study judged that the modular HTGR concept with prismatic-block fuel and a steam-cycle power conversion system was ready for commercial demonstration in the United States in the “near future,” based on past demonstrations that it deemed successful, the billions of dollars already spent on the technology, and foreign construction projects. The DOE argued that these designs could be commercially available sometime in the 2030s. In contrast, less mature designs, which would first require engineering and performance demonstrations, would not be commercially available until around 2050 (Petti et al. 2017).

While the DOE study is silent on whether it considers pebble-bed designs to be of comparable maturity to prismatic-block designs to support near-term commercial demonstration, in October 2020 the agency chose X-Energy's Xe-100 pebble-bed four-pack as one of two commercial demonstration projects to be built by 2027 under the ARDP.

This timeline is universally acknowledged to be aggressive. In 2018, X-Energy developed a schedule that would have its first commercial pebble-bed unit operating by the early 2030s, assuming about four years for long-lead procurement and five years for reactor construction (Bowers 2018). But soon afterward, it proposed an accelerated development program leading to an “efficient and effective commercial demonstration” by the mid- to late 2020s, consistent with the ARDP timeline (DOE-NE 2019).

However, as is the case with sodium-cooled fast reactors, it is far from clear that the past U.S. HTGR projects have

provided a sufficient technical basis for skipping the performance demonstration step. As discussed earlier, the US demonstrations at Peach Bottom and Fort St. Vrain had a “mixed” track record, according to the DOE. And those reactors had significant differences from the current generation of modular prismatic-block HTGR designs—especially when it comes to safety.

PERFORMANCE DEMONSTRATION: THE NECESSITY OF PROTOTYPE TESTING

The safety basis for current HTGR designs relies on the effectiveness of passive cooling and the robust performance of the fuel. However, past HTGR demonstrations did not employ passive cooling systems—in fact, the 842 MWth Fort St. Vrain reactor was too large to have had a passive cooling system. Also, these demonstrations used fuels that were significantly different from current-generation fuels—for example, they contained HEU and thorium, rather than HALEU, and they used uranium carbide or oxide kernels, rather than the uranium oxycarbide that would be used for the Xe-100 and other designs today. There has never been an integrated demonstration of the safety features of a small modular HTGR that could support claims that the reactor does not need a robust containment or off-site emergency planning measures.

This is not a new issue. In the 1980s, the DOE proposed building a commercial demonstration plant based on a modular prismatic block HTGR design without a conventional containment at an unspecified but “typical” nuclear power plant site (Williams, King, and Wilson 1989). (This was in contrast to the DOE's plan around the same time to build a prototype of the PRISM fast reactor, as discussed in chapter 5.) In its (draft) “pre-application” safety review, the NRC staff rejected this approach, concluding that:

based on judgments of the adequacy of existing operating experience, the novel design features proposed, and the status of the present technology base, the staff requires that testing and operation of a prototype test reactor, located at an isolated site, be mandatory before design certification (Williams, King, and Wilson 1989).

In NRC parlance, a “prototype” is a “nuclear reactor . . . used to test design features or new safety features” and “can be . . . a standard plant design in all features and size, but may include additional safety features to protect the public and the plant staff from the possible consequences of accidents during the testing period” (NRC 2017). Thus, an NRC prototype would be comparable to the DOE's concept of a performance demonstration reactor, and could have additional safety features not included in the commercial version. (An NRC

prototype could itself serve as a commercial reactor if it were licensed under the appropriate provisions, but the additional safety requirements and stipulation of an isolated site might not be attractive features for a commercial customer.)

The NRC staff argued in the draft safety review that such prototype testing was necessary “to compensate for removal of the traditional (and testable) containment building” and to “help ensure that licensed plants of that design have adequate fission-product retention” (Williams, King, and Wilson 1989). But the DOE project did not go forward, and neither a prototype nor a commercial demonstration plant were built.

Today, the DOE is pursuing a commercial HTGR demonstration plant—without a traditional containment—without first conducting the prototype testing that the NRC staff had previously said was essential. In recent decades the NRC has weakened its policy on the need for prototypes in NLWR licensing, and it is unclear whether the staff would take the same position today as it did in the 1989 HTGR draft safety review (NRC 2017). However, the concerns that the NRC staff raised at that time are still valid today, and the agency should require prototype testing given the outstanding uncertainties in the HTGR safety approach discussed above. Even without an NRC requirement to do so, it would be highly prudent for any vendor to demonstrate the safety features that would justify the absence of a containment.

The Department of Defense (DOD) is proceeding more cautiously than the DOE in pursuing the development of mobile microreactors for deployment at military bases. In March 2020 DOE awarded X-Energy \$14.3 million to first develop a prototype design for its 10 MWth mobile microreactor concept before deciding whether to move forward.²²

Prototype testing will be even more important for pebble-bed designs, which—unlike prismatic-block HTGRs—have never been demonstrated in the United States. The German experience with pebble-bed reactors was mixed. And the only currently operating pebble-bed reactor, the HTR-10, cannot be considered a performance demonstration or even a full engineering demonstration. It has operated only intermittently and has not demonstrated key pebble-bed systems and safety features. The reactor only ran at full power for a brief period, and major features of the all-important system for on-line fuel loading/discharge/reloading had not been used by 2015, as the reactor core had still not achieved equilibrium (steady-state) operation at that time (Knol et al. 2015).

In addition, safety tests that had been planned at the HTR-10 in a collaborative program with the European Commission, including a melt-wire test to validate temperature calculations, were never carried out (Knol et al. 2015). This was due in part to China’s preoccupation with designing and building the HTR-PM commercial demonstration (Knol et al.

2015). In retrospect, China’s decision to proceed with the now-delayed HTR-PM without first undertaking a full complement of testing and demonstrating reliable equilibrium operation at the HTR-10 is probably a mistake that the United States should not repeat.

FUEL QUALIFICATION, SAFETY TESTING, AND COMMERCIAL-SCALE FABRICATION

The other key HTGR safety feature, high-integrity TRISO fuel, also will require further development and qualification prior to a commercial reactor demonstration. And of course facilities will need to be built to produce the fuel for the reactor.

FUEL QUALIFICATION AND TESTING

Given that a typical program for new fuel qualification can take 15 to 25 years, this could well be the rate-limiting step for new reactor development. To arrive at the optimistic, 13- to 15-year deployment schedules for a commercial demonstration reactor, vendors such as X-Energy have assumed that fuel qualification; fuel manufacturing capability; and reactor design, licensing, and construction can all occur concurrently, rather than sequentially (Bowers 2018). However, this approach may not be adequate to ensure that the fuel will perform as well as advertised.

The HTGR does have a fuel development advantage compared to other NLWRs given that the DOE has been conducting the Advanced Gas Reactor fuel qualification program described above since 2002. However, the schedule for completing the program, including the critical post-irradiation examination of the fuel, has slipped at least five years from its original 2020 completion date. Moreover, due to the problems encountered during the irradiations described above, there are gaps in the data and the program may never fully achieve its goals. In particular, the AGR-5/6/7 irradiation tests, which were intended to serve as a formal program to qualify TRISO fuel under both normal and abnormal operating conditions, were terminated prematurely (Pham et al. 2020). It is unclear whether the data that was collected will be adequate for finalizing fuel fabrication specifications and for the requisite NRC approval for use of the fuel.

In 2019 the Electric Power Research Institute (EPRI), the main US nuclear industry research organization, asked the NRC to issue a finding that “testing of UCO TRISO-coated fuel particles in AGR-1 and AGR-2 constitutes a performance demonstration of these particle designs over a range of normal operating and off-normal accident conditions” (EPRI 2019). It is not entirely clear what that would mean in a regulatory context—as the NRC pointed out, “...how the TRISO fuel meets regulations will depend on how the design and other systems, structures, and components are credited in the overall

safety of the [reactor] design” (NRC 2020a). It does appear to be a suggestion that the NRC approve the use of TRISO fuel without having the benefit of the AGR-5/6/7 qualification program data or of post-irradiation accident testing, including critical air and water exposure tests. However, the AGR-1/2 tests alone provided little useful data for TRISO fuel qualification. The AGR-1 test used fuel that was carefully fabricated under laboratory conditions and is of limited relevance for fuel produced under commercial conditions. Moreover, it is premature to reference the AGR-2 test, since AGR-2 fuel particles are still undergoing post-irradiation examination and final results are not yet available. And the thermocouple failures throughout the experimental campaign made it impossible to collect critical irradiation temperature data. Ultimately, the NRC approved a far more limited statement on the utility of the AGR-1/2 test data than EPRI had requested.

Even data from the entire AGR test series likely will be insufficient to complete HTGR fuel qualification. The tests were conducted in the Advanced Test Reactor, which is water-cooled and, as such, does not fully replicate the conditions within an actual HTGR, such as the neutron energy spectrum. Also, the tests did not use representative fuel compacts for any current HTGR design (either prismatic-block or pebble-bed), so the data are only valid for the TRISO particles themselves and not for entire fuel assemblies (Sunseri 2020). Finally, the test irradiations were conducted in an accelerated manner by subjecting the fuel to a higher neutron flux than it would be exposed to in a reactor. While these may be adequate for the initial stages of fuel development, an HTGR prototype would be needed to test the fuel under more realistic conditions.

In a past review of the AGR program in the context of the now-cancelled Next Generation Nuclear Plant (NGNP) project, the NRC staff stated that “the lack of fuel performance data obtained in real-time HTGR neutron environments is of concern to the NRC staff. This concern is based on the questionable adequacy of data generated solely in accelerated irradiation environments” (NRC 2012), observing that “reliance on NGNP prototype testing may be necessary to adequately demonstrate design safety features associated with fuel, core, and reactor system performance.” But the DOE did not build a prototype demonstration reactor project that it could have used to optimize fuel design and complete TRISO fuel qualification. Instead, the DOE is proceeding directly with construction of a commercial demonstration HTGR—the Xe 100.

CONSTRUCTION OF INDUSTRIAL-SCALE FUEL FABRICATION FACILITY

In addition to the reactor itself, a commercial HTGR demonstration project such as the Xe-100 will require the design,

development, licensing, and construction of an industrial-scale fuel fabrication facility—itsself a time-consuming and costly undertaking.

The prospects for these facilities in the United States are uncertain. The United States does not currently have the capability to fabricate TRISO fuel on a commercial scale. To meet its timeline for reactor deployment, X-Energy intends to submit an application for a TRISO fuel fabrication facility (called TRISO-X) by early 2021. And BWXT is moving forward with a plan to restart the TRISO fuel manufacturing line at its facility in Lynchburg, Virginia, that was used to produce some of the fuel for the AGR irradiation program (although with inferior quality, as discussed earlier).

Although X-Energy originally intended to locate TRISO-X in an existing building at Oak Ridge National Laboratory, it now apparently plans to install it within the General Electric-Hitachi Global Nuclear Fuel uranium fuel fabrication facility in Wilmington, North Carolina. However, as discussed in chapter 4, the security, material accounting, and criticality safety programs at that facility will have to be upgraded to meet Category II standards, since it is currently only licensed to handle Category III uranium with an enrichment below 5 percent. These modifications could prove disruptive to the facility’s operations (NEI 2018).

The facility may also have to produce different types of TRISO fuel at the same time, which will require greater capacity and complicate operations. X-Energy now has both DOD funding that may lead to an Xe-Mobile microreactor prototype and DOE funding for an Xe-100 four-pack commercial demonstration project—both to be constructed over the next several years. X-Energy President Harlan Bowers said in early 2019 that the company would “‘gear the [TRISO-X] design to serve’ whichever reactor is first to the market” (Freebairn 2019).

In any case, X-Energy’s aggressive schedule would need the TRISO-X facility to be licensed and operational by 2023–2024, before the final results of the AGR fuel development and qualification program will have been obtained. However, if those results indicate that modifications to the fabrication process are needed, the production plant may need to be retrofitted. Given the complex interactions between fuel performance and reactor design and operation, there is risk in moving forward with commercial-scale fuel fabrication licensing and construction before fuels have been fully developed and qualified through prototype testing. Consequently, it would be prudent to postpone fuel fabrication facility design until an advanced stage of fuel development has been achieved and the necessary process parameters finalized.

HALEU AVAILABILITY

A related obstacle to near-term deployment of prototypes and commercial demonstration HTGRs is the availability of the considerable quantities of HALEU that the reactors will need. As with some other NLWR designs, HTGRs will require a source of HALEU at an enrichment level not currently produced at commercial uranium enrichment plants.

Each of X-Energy's Xe-100 200 MWth modules, for example, would require about 1.5 metric tons heavy metal of 15.5 percent-enriched HALEU for the initial core (equivalent to the current worldwide demand for US HALEU), and about 0.45 metric ton heavy metal of fresh fuel would be required annually. Even a 10 MWth TRISO-fueled microreactor would require a few hundred kilograms of HALEU. As discussed in chapter 4, the Nuclear Energy Institute has estimated that it would take a minimum of seven to nine years to establish the fuel cycle infrastructure to support a significant level of HALEU production, assuming full funding is available (NEI 2018). The only current prospect for enriching HALEU is the Centrus Energy Corporation's three-year, DOE-funded pilot centrifuge demonstration project in Piketon, Ohio, but that will produce, at most, multi-kilogram quantities by 2021, and only up to 900 kilograms per year at peak capacity (Dyke 2020).

COST AND FINANCING

In addition to the timeline for fuel availability, there is the all-important question of demonstration project cost and available financing, about which public information is scarce. The maximum amount of funding that the DOE ARDP has committed to X-Energy for its four-unit demonstration plant (subject to availability of future congressional appropriations) is \$1.6 billion, provided X-Energy can match that amount. But it is not clear the total amount of \$3.2 billion will be sufficient for the capital cost of a four-reactor plant. While X-Energy has claimed that a single 200 MWth reactor plant would cost less than \$1 billion (Bowers 2017), DOE researchers estimated a 200 MWth prismatic core test reactor would have a capital cost around \$2 billion to \$6 billion (2020 dollars), with a best estimate of around \$4 billion (Sterbentz et al. 2016). There is also the capital cost of the TRISO fuel fabrication plant, which X-Energy has estimated at \$100 million to \$200 million, two-thirds of which the company hoped to finance by debt (Freebairn 2019).

And there are the operating costs, which include the cost of fuel, operations, and maintenance—for which X-Energy presumably would be fully responsible. For a 200 MWth Xe-100 module, the annual cost of the required HALEU would be at least \$2.5 million today (assuming the separative work unit cost to produce HALEU would be equal to that available on the open market, which is questionable given there is no current supply). DOE researchers estimated

TRISO fuel fabrication for a first-of-a-kind plant would cost \$26,500 per kilogram of uranium in 2009 dollars (or around \$32,000/kg in 2020 dollars) (INL 2012). At this rate, fuel would cost nearly \$18 million per year for each Xe-100 module (it is not clear whether this estimate includes the annualized fabrication facility capital cost). And the researchers estimated the operating cost of a 200 MWth test reactor at around \$20 million to \$60 million per year, with a best estimate around \$40 million per year (2020 dollars). The best-estimate annual operating cost alone, not including any financing payments, would be \$58 million per year, or \$88 per MWe-hour, nearly three times the average generating cost of the current reactor fleet of less than \$31 per MWe-hour, illustrating the economic hurdles faced by this technology.

Finally, there is the question of whether X-Energy has the requisite funding to provide a 50-50 match of the DOE contribution. In 2018, it was reported that X-Energy had “nearly \$39 million in private investment” (Walton 2018), and the company received prior funding commitments from the DOE totaling less than \$60 million. Unless the company can realize its goal of a capital cost of less than \$1 billion for an Xe-100 module, it will need to raise far more money. Another track for a demonstration HTGR is the funding provided by the DOD to X-Energy and other vendors for a defense micro-reactor prototype, but that project is not likely to lead to development of economical power reactors for civilian use.

Although the exact cost of an HTGR demonstration is uncertain, what is clear is that graphite-moderated HTGRs will be “costly to build and operate” because of their large size and low power density (Duchnowski et al. 2019). To address this cost penalty, in 2019 the DOE awarded grants of \$2 million to Stony Brook University for developing alternative moderator materials utilizing hydrogen or beryllium, and \$3.5 million to X-Energy to examine ways to “reduce construction and maintenance costs” of the Xe-100 reactor design (DOE-NE 2019).

In summary, HTGRs do have some attractive safety features, but they also have a number of drawbacks. And the push by the DOE and HTGR designers to reduce defense-in-depth features such as containment strength and emergency planning zone size could undermine any safety benefits the design offers. Moreover, the reactors do not appear to meet the DOE's sustainability goals, given their unfavorable characteristics with regard to uranium utilization and waste generation. Overall, it is difficult to assess whether HTGRs would represent an improvement over LWRs. Given the mixed performance of previous HTGR demonstration reactors and fuels, as well as their differences from current designs, it is likely that additional engineering and full-scale demonstrations will be needed to resolve outstanding safety and performance issues before HTGRs could be deployed on a commercial scale.

Molten Salt Reactors

Just about anyone with an interest in nuclear energy and an Internet connection has probably heard about molten salt reactors (MSRs). MSRs differ from other reactor types in that they use a hot, liquified salt (which can include such compounds as sodium chloride, or table salt) to cool the reactor and transfer heat to a power conversion system to generate electricity. In most MSR designs the nuclear fuel itself is dissolved in the molten salt coolant and is thus in a liquid form. These types of MSRs are significantly different from conventional LWRs or other NLWRs that use solid fuels.

MSRs have been promoted for years by enthusiastic advocates as reactors with major cost, safety, nonproliferation, and sustainability advantages over LWRs or other types of advanced reactors (Pellum 2019). One popular MSR design called the Liquid Fluoride Thorium Reactor (LFTR (pronounced “lifter”)), would use a fuel salt containing the element thorium to breed the fissile isotope uranium-233 (U-233), a process that in theory could be self-sustaining, similar to a fast plutonium breeder reactor. The potential for development of a thermal-spectrum breeder reactor was one of the original motivations for pursuing the MSR in the 1960s.

In theory, MSRs are very flexible. They can use solid as well as liquid fuel. They can use thermal (slow) neutrons, fast neutrons, or intermediate-energy (“epithermal”) ones. They can be fueled with mixtures of thorium and U-233, highly enriched uranium (HEU), low-enriched uranium (LEU), or even transuranic elements (TRU) such as plutonium extracted from spent nuclear fuel. They can use one fluid loop or multiple loops. They can operate as breeders or burners. However, as with any reactor concept, in practice

there are engineering and safety limits that greatly narrow the range of workable designs.

In most liquid-fueled MSRs, the flowing nuclear fuel also serves as the primary coolant system, transferring heat out of the reactor core to a secondary coolant system. The secondary coolant system also typically uses a (non-fuel) molten salt to transfer the heat generated by the reactor to an electricity conversion system, which would likely be a conventional steam turbine. (Although some MSR designs, like HTGRs, could potentially use an advanced helium-gas-powered turbine, that technology remains relatively immature and is not likely to be available in the near term.)

MSRs can be operated at a lower pressure than LWRs, but must be constantly maintained at a high temperature in order to keep the salt in a liquid state. The required salt temperature ranges from 650°C to 750°C under normal conditions. One drawback is that molten salts are highly corrosive to many structural materials. A key technical challenge in building MSRs is to find materials that can tolerate prolonged exposure to the fuel at high temperatures in high radiation fields.

The main advantage of liquid reactor fuel is its potential to achieve a higher fuel burnup and higher conversion ratio than solid fuels, which could reduce nuclear waste volume and increase uranium utilization. (The burnup is related to the fraction of the initial heavy metal content of fuel that is converted to heat energy, and the conversion ratio is related to the ability of the reactor to convert fertile materials such as U-238 or thorium-232 to fissile materials.) As discussed in chapter 3, the fuel burnup in solid-fueled reactors is constrained by changes in safety and performance that occur as fuel is irradiated in a reactor. These include physical changes in the structure of the crystalline fuel and cladding,

and changes in the fuel composition as uranium and other fuel materials are bombarded by neutrons and transmuted into other elements.

In contrast, liquid fuels lack the solid structure that can be damaged by irradiation, and the fuel composition can be adjusted as needed to optimize its performance simply by blending in additional liquid streams. However, compared to solid fuels, liquid nuclear fuels introduce numerous additional safety, environmental, and proliferation risks. In solid-fueled nuclear reactors, the fission products generated during reactor operation are largely trapped by the molecular structure of the fuel pellets or the cladding that surrounds them. Some fission products, such as the noble gases krypton and xenon, migrate to the surface of the fuel pellets, but they remain confined within the fuel pins as long as the cladding is intact. In molten salt–fueled reactors, these fission products are released from the reactor core and must be either retained by the reactor structures and off-gas treatment system or released to the environment.

This has implications both for safety and for nuclear nonproliferation. In particular, the very large releases of noble gas fission products from MSR designs could interfere with the functioning of the international monitoring system established under the Comprehensive Nuclear Test Ban Treaty to detect the occurrence of clandestine nuclear weapon tests. In addition, most liquid-fueled MSR designs require continuous reprocessing of the fuel salt in order to adjust the fissile material content and remove fission products trapped in the salt that may reduce reactor performance and safety. As discussed below, continuous reprocessing will pose unique difficulties for application of International Atomic Energy Agency (IAEA) safeguards and will increase proliferation risks. These risks must be fully addressed and mitigated if MSR designs are to play a significant role in the future energy mix.

History and Current Status

In comparison to liquid metal–cooled fast reactors and gas-cooled reactors, there is much less operating experience with thermal MSR designs and none at all with fast MSR designs. There have only been two engineering-scale demonstrations. The 1940s Aircraft Reactor Experiment at Oak Ridge National Laboratory was a small experimental test reactor that operated for only nine days. The subsequent Molten Salt Reactor Experiment (MSRE), a 7.5 megawatt-thermal (MWth) test reactor, operated at Oak Ridge from June 1965 to December 1969.

MOLTEN SALT REACTOR EXPERIMENT (MSRE)

The MSRE operated only intermittently over its four-year lifetime and was critical for a total of about two years of that

time span. The average power achieved over that time was about 4.18 MWth, including shutdown periods, in comparison to its maximum power of about 7.5 MWth. The fuel was a mixture of lithium fluoride, beryllium fluoride, and uranium fluoride. The initial uranium fuel was HEU with a U-235 enrichment of about 33 percent; later it was changed to U-233. (The U-233 was supplied from other reactors and was not generated at the MSRE itself.) A small quantity of plutonium-239 was added later for good measure.

After the MSRE was shut down, the Atomic Energy Commission concluded that the work to date had not “advanced the program beyond the initial phase of research and development” and that “about 2 billion dollars in undiscounted direct costs [more than \$12 billion in 2020 dollars] could be required to bring the molten salt breeder . . . to fruition as a viable, commercial power reactor” (AEC 1972). More recently, a Department of Energy (DOE) review concluded that the MSRE “should be considered test laboratory scale or perhaps engineering scale” and pointed out that it was never connected to a power conversion system (Petti et al. 2017).

RECENT DEVELOPMENTS

Numerous small companies have started up in the last few years to pursue development of various MSR designs. In the United States, companies including Terrestrial Energy, ThorCon, and Flibe Energy are pursuing liquid-fueled thermal MSR designs. (Another company that was developing an MSR, Transatomic Power, went out of business.) TerraPower and Elysium Industries are pursuing liquid-fueled fast-spectrum MSR designs. Moltex is developing a fast reactor that would use stationary, metal-clad fuel elements containing molten salt fuel instead of solid fuel. Also, Kairos Power is developing a molten fluoride-salt-cooled, high-temperature reactor that uses a solid fuel similar to TRISO pebble-bed HTGR fuel.

The DOE itself has been slow to get on the MSR bandwagon. The MSR was one of the few concepts that the DOE did not initially fund in its Generation IV program. More recently, responding to growing private sector interest, the DOE began providing a modest amount of support for MSR research. In 2016, it committed to providing up to \$40 million over five years to a consortium including Southern Company and TerraPower for basic research and development on a fast molten chloride salt reactor—apparently the first US government funding for a liquid-fueled MSR project in 40 years. Several other grants to other companies and researchers have followed, including (now-defunct) Transatomic Power, Terrestrial Energy, ThorCon, and Flibe Energy, ranging from several hundred thousand to a few million dollars each.

In December 2020, the DOE provided \$30 million in initial funding to two MSR designs as part of a second-tier of “risk

reduction” awards under the Advanced Reactor Demonstration Program (ARDP). One award went to Kairos Power to develop a reduced-scale test reactor called Hermes based on its molten salt–cooled, TRISO-fueled design, and the other to Southern Company for the “Molten Chloride Reactor Experiment”—also presumably a reduced-scale test reactor. Kairos ultimately will receive \$303 million from the DOE of the estimated \$629 million cost of the Hermes, and Southern up to \$90.4 million of the projected \$113 million cost of its test reactor (DOE-NE 2020).

Many of the liquid-fueled MSR startup companies have emphasized sustainability in touting the benefits of their designs. For example, Elysium Industries advertises that its fast-neutron MSR “has the ability to consume spent nuclear fuel and weapons waste transforming it into useful energy” (Elysium Industries n.d.). Indeed, the operational flexibility of liquid fuels could potentially allow for significant improvements in sustainability compared to solid-fueled reactors. However, as discussed below, these claims are exaggerated, and in at least one case (the defunct company Transatomic Power) demonstrably false.

Safety

The use of liquid fuel instead of a solid fuel in an MSR has significant safety implications for both normal operation and accidents.

NORMAL OPERATION

In contrast to solid-fueled reactors, molten salt–fueled reactors release a large fraction of the gaseous fission products produced in their liquid fuel after they are generated. Since the gases do not escape from the fluid rapidly enough on their own, they must be constantly removed to avoid increasing the fluid pressure and decreasing the reactivity of the fuel. This is done by circulating a stream of helium gas bubbles to push the fission product gases from the fuel. The fission product gases must then be trapped and either (1) stored for a long period of time (years to decades) and eventually disposed of in a geologic repository, or (2) stored for a short period of time to allow for some radioactive decay and then released into the environment.

Safe management of these fission products will be a much greater challenge for MSRs than for LWRs (Lyman 2019). LWRs do release some noble gas fission products during normal operation because a small number of fuel rods will experience cladding failures. However, the noble gas releases from a MSR core would be hundreds of thousands of times greater than the releases from LWRs of comparable capacity (Lyman 2019). In principle, short-lived fission products, such

as xenon-135 (with a half-life of 9.1 hours), can be captured and stored until they have decayed away. Other, longer lived noble gas fission product isotopes, such as krypton-85 (with a half-life of more than 10 years), are more challenging to manage. But in either case, it is difficult and expensive to capture and store large flows of noble gases, because they are chemically inert. Operators of liquid-fuel MSRs may face significant challenges in meeting safety limits on reactor discharges of noble gas fission products into the environment.

The human health and environmental impacts of chemically inert noble gas fission products such as krypton-85, although significant, are relatively low compared to isotopes such as cesium-137. Noble gases disperse quickly into the atmosphere when released. However, cesium-137 deposits on and binds to soil and other surfaces, and thus is one of the most environmentally hazardous radionuclides. With a 30-year half-life, cesium-137 is largely responsible for the persistent radiological contamination in the regions around Chernobyl and Fukushima.

MSR developers often say that cesium-137 and certain other troublesome fission products do not present a problem because they remain chemically bound in the liquid fuel and are not released. However, this is misleading because it only applies to those isotopes that are generated directly in the fuel from fission. It does not apply to isotopes that are produced indirectly by the decay of noble gas fission products only *after* the gases are released from the fuel.

For instance, in addition to being produced directly by fission, cesium-137 also results from the decay of short-lived xenon-137, which has a 3.82-minute half-life. Indeed, nearly all of the cesium-137 generated in a nuclear reactor is produced through xenon-137 decay rather than directly by fission. In a solid-fueled reactor, most of this cesium remains trapped in the fuel unless the fuel is damaged during an accident. But in an MSR, almost all of the cesium-137 generated is released from the fuel under normal conditions, and it must be captured and safely stored. In a two-month period of normal operation, a 1000 megawatt-electric (MWe) MSR could release about as much cesium-137 from the core as the total amount released into the environment from the Fukushima accident. Indeed, MSRs are such good generators of cesium-137 that Oak Ridge scientists received a patent in 1972 for a method for production of “high-purity cesium 137” utilizing the MSRE off-gas stream (Lyman 2019).

One type of MSR, Moltex, was designed in part to address the problem of cesium-137 release. As mentioned above, Moltex uses a molten salt fuel that is not free-flowing but is contained in metal-clad fuel rods. This design allow the fuel cladding to trap xenon-137 long enough for it to decay to cesium-137, although at the expense of losing the flexibility

to easily adjust fuel compositions (Scott 2017), one of the major motivations for MSR. (Moltex fuel rods would be periodically vented to release longer-lived noble gas isotopes, however.) But other MSR developers say little about this issue and how they expect reactor operators to manage and dispose of large quantities of high-level cesium-137 waste.

Other fission product impurities—in particular, noble metals such as ruthenium—are not soluble in the molten salt and also must be continuously filtered out. Otherwise, they would collect on metal structures in the reactor, such as the heat exchangers, creating hot spots that could damage the structures (Forsberg 2006).

Another troublesome radionuclide, tritium, with a half-life of 12.5 years, is highly mobile and cannot be effectively captured. Even with a costly off-gas control system, MSRs would almost inevitably discharge far more tritium and other radioisotopes into the environment during normal operation than solid-fueled reactors.

ACCIDENTS

Advocates for liquid-fueled MSRs cite three main benefits of MSRs compared to LWRs that they claim would reduce the risk of accidents (LeBlanc 2010). First, they note that, in contrast to solid fuel, the liquid fuel cannot melt down in an accident. Second, they state that the reactors are passively safe because (for some designs) if the molten fuel overheats, it would quickly drain into a place where it could be safely cooled and stored. Third (and similar to an argument made for sodium-cooled fast reactors), they point out that since the reactor operates near atmospheric pressure, there is less risk of a leak of radiation to the environment and the reactor structure does not have to be qualified to withstand high pressures. Each of these points is considered below.

LIQUID VS. SOLID FUEL

The observation that the core of a liquid-fueled reactor cannot “melt down” is a tautological statement that gives a misleading impression of the safety of such reactors. In the event of a severe accident or sabotage, the timing and size of a release of radioactive material into the environment are key considerations. The fact that the fuel in an MSR is already a liquid does not confer a safety advantage in these respects.

In an LWR, fission products are largely trapped within the molecular structure of the solid ceramic fuel pellets or the metal cladding that surrounds them. If cooling is lost, the fuel will heat up, the cladding will become damaged and rupture, and the fuel pellets will eventually begin to melt. Some fission products will become mobile as the fuel softens and may escape from the fuel. Eventually, the highly corrosive molten fuel will drop to the bottom of the reactor vessel and will melt

through it, spilling onto the containment floor. At that point, the containment is the only remaining barrier to release of radioactive material into the environment.

In the most severe LWR accidents, cladding damage and the release of some fission products from the fuel into the coolant could begin as soon as 30 minutes after cooling is lost. However, because of the high melting point of the uranium dioxide ceramic fuel used in LWRs (around 2800°C), it typically takes several hours until the fuel starts to melt, and many more hours until the molten core breaches the bottom of the reactor vessel and flows into the containment. This provides time to take measures to mitigate the accident, such as restoring cooling, and to implement emergency plans to protect the public.

After core damage occurs in an LWR accident, the rate at which radioactive material is released to the environment depends on the extent to which remaining barriers, such as the reactor vessel, piping, and containment structure, are still capable of retarding releases. Characterizing the actual source term (the timing and composition of fission product releases) for any accident scenario is quite complex and involves substantial experimental and analytical work. For example, during the 2011 Fukushima accident, damage to the core of the first reactor did not begin until about three hours after the reactor lost all cooling, complete melting of the fuel took several more hours, and radiation releases to the environment did not begin until about 12 hours after the loss of cooling. But as bad as the accident was, the containment shells at the three damaged reactors remained largely intact, and the overall releases of highly hazardous fission products such as iodine-131 and cesium-137 were only a few percent of the total amounts that could have been released.

In contrast, in an MSR there is no fuel cladding, and the fuel itself is initially in a highly corrosive liquid state instead of a stable solid. The reactor structures containing the fuel, rather than the fuel rods themselves, form the first line of defense for fission product release. As discussed above, even during normal operation, the fuel continuously releases gaseous fission products that are either captured or eventually released. In the event of an accident, the fuel could heat up rapidly to the point where it would start to release additional radionuclides at a much higher rate. However, very little information is available about the physical properties of molten salts that would shed light on the radiological source term of such an event. Molten salts also must be maintained at a high temperature (over 600°C) to remain in the liquid state; if areas within the fuel get too cold, the salt can crystallize and clog pipes, blocking coolant flow and ultimately resulting in a dangerous temperature increase (IRSN 2015).

After radionuclides escape from the fuel, releases to the environment would depend on the integrity of the remaining

barriers—namely, the reactor structures and the off-gas treatment system. But the structural materials currently available would have a very limited capacity to contain the molten fuel at the very high temperatures that could occur during an accident. For that reason, the over-heated fuel must be cooled very rapidly—for instance, by draining it from the core into a special chamber—or the reactor could be destroyed, releasing large amounts of radioactivity into the environment.

PASSIVE SAFETY

To mitigate a severe loss of cooling, most MSR designs include a safety feature called a freeze plug. One or more plugs of frozen salt are used to close off a drain at the bottom of the reactor vessel. In the event of a loss of cooling or loss of external power causing fuel overheating, the plugs would melt before the fuel reached a dangerous temperature, allowing the fuel to drain quickly into dump tanks below the reactor vessel. The dump tanks would be designed to maintain the discarded fuel at a safe temperature and in a configuration where it could not become critical and start generating power again.²³

Operation of the freeze plug sounds simple in theory, but is far more complex in practice. For instance, it is not clear whether the local decay heat of the fuel would be sufficient to rapidly melt the freeze plugs, or whether an external heating source would be needed (in which case the mechanism may not be entirely passive and would not function if external power were lost). Also, to judge the effectiveness of this safety mechanism, one must calculate how long it would take for the plugs to melt and the fuel to completely drain. The core would have to drain quickly enough to avoid destroying the reactor structures that contain it.

The few studies of these issues to date have shown that MSRs will heat up rapidly in the event that cooling is interrupted, leaving very little time to mitigate the accident if the fuel fails to drain. A 2013 study of the fast-spectrum Molten Salt Fast Reactor (MSFR) being designed in France found that in the event of a station blackout or other accident causing a loss of heat removal, it would take as little as eight minutes for the core to heat up to 1200°C, the temperature at which the structural materials are assumed to fail (Brovchenko et al. 2013). Other studies have estimated grace periods of up to 22 minutes for this reactor, but researchers point out that “the MSFR design has not been finished, and no detailed thermohydraulic studies have been conducted which would give accurate information specific to the kind of accident scenario expected to trigger the melting of the freeze plug” (Shafer 2018). Thus, there is “no definitive estimate” of the time it would take for an MSR to heat up to 1200°C (Tiberga et al. 2019).

Therefore, should such an accident occur, only tens of minutes at most may be available for the freeze plugs to melt

and the fuel to drain completely from the core to avoid a structural collapse and large radiological release. It remains unclear whether this is achievable in practice. One study has shown that this can be accomplished in as little as 95 seconds; however, if the freeze plugs only partially melt or are blocked by solidified fuel, the drain time could be increased from 95 seconds to more than 20 minutes (Wang et al. 2016). A more recent study concludes that “a freeze-plug design based only on the decay heat to melt is likely to be unfeasible” (Tiberga et al. 2019). Given the complexity of the system, uncertainties are large, but—given the short timelines—there is very little room for error.

Some MSR designers are not taking the passive freeze plug for granted. The French MSFR design includes both active and passive drain valves. In addition, the design deliberately introduces a weak spot in the floor of the cavity containing the reactor. The idea is that this area will preferentially fail near the drain valves, so that the fuel will be funneled into the dump tanks in the event of a catastrophic failure (IRSN 2015). Terrestrial Energy has eliminated freeze plugs from its thermal MSR design; instead, it limits core size with the expectation that natural convection cooling would be sufficient to prevent the core from heating up to a dangerous temperature.

However, there are questions about whether passive cooling methods alone would be adequate. According to the DOE, “the full range of design basis accident scenarios [for MSRs] has not been established, so the need for active safety systems cannot be ruled out” (Petti et al. 2017).

REACTIVITY FEEDBACK

Another passive safety feature is inherent negative reactivity feedback—the tendency for the nuclear chain reaction to shut down if the reactor heats up. As discussed in chapter 2, LWRs have this property. In MSRs, however, the feedback behavior turns out to be very complex. While the thermal spectrum MSRE developed by Oak Ridge National Laboratory, with a graphite moderator, was originally thought to have negative reactivity feedback, this was discovered to be incorrect when the system was analyzed decades later with more modern and accurate methods (Mathieu et al. 2006). Other thermal spectrum MSRs, such as the one designed by Transatomic Power, have very complex reactivity behavior, and at certain times during the operating cycle will have positive moderator or void coefficients (Robertson et al. 2017).

Partly as a result of the finding that thermal MSRs can have positive reactivity feedback, European researchers decided to pursue fast-spectrum MSRs, which have no moderator, such as the French-designed MSFR discussed above (Mathieu et al. 2009). Unlike their solid-fueled fast

reactor cousins, fast MSR typically have negative reactivity feedback coefficients, making them more stable. However, fast MSR have other safety drawbacks, including a large fuel volume that would have to be rapidly drained in the event of a loss of cooling.

LOW PRESSURE

Although MSR operate at higher temperatures than LWRs, they operate at lower pressures, which could be advantageous for safety. According to the International Atomic Energy Agency (IAEA), “low operating pressures can reduce the risk of a large break and loss of coolant as a result of an accident, thereby enhancing the safety of the reactor” (IAEA n.d.). However, as discussed in chapter 1, accidents in which temperatures and pressures quickly rise would also be a concern. The low system pressure also introduces risk because water can more easily flow into the reactor, which could cause a violent steam explosion (IRSN 2015). And as with the HTGRs discussed in chapter 6, the water could also react with the graphite in thermal MSR, such as Terrestrial Energy’s Integral Molten Salt Reactor (IMSR). If the reactor were flooded as the result of a natural disaster, the overpressure could force water to leak into the reactor. And a terrorist group could sabotage an MSR simply by pumping water into the core.

ADDITIONAL SAFETY CONCERNS

Other safety concerns arise with the co-located, on-line waste processing facilities and/or off-gas treatment systems that MSR would require. Having chemical processing operations in proximity to an operating reactor would introduce an entire set of accident scenarios not encountered at LWRs. In addition, the large quantities of fission products that would be stored outside of the core in off-gas treatment systems, including both noble gases and their decay products, such as cesium-137, would pose additional risks. An accident or terrorist attack on those waste processing facilities could be as severe as one affecting the reactor itself—or even worse.

THE BOTTOM LINE

Compared to LWRs, MSR offer a number of safety disadvantages and only minor safety benefits. On the negative side, the liquid form of the fuel allows for far greater releases of radiation from the core under normal conditions and more rapid releases to the environment during accidents. The MSR has fewer levels of protection against fission product release than does an LWR. In particular, there would be additional risk from the storage of fission products such as cesium-137 in the radioactive waste generated by the off-gas system.

On the positive side, the low pressures at which MSR operate may reduce the likelihood of pipe ruptures resulting in large-break loss of coolant accidents. However, the high

temperature and power density of molten salt fuel could also increase the risk of other types of accidents that would affect cooling. For an MSR, preventing the release of radiation in such scenarios is largely dependent on being able to drain the overheating fuel in a matter of minutes to avoid a structural collapse and fuel vaporization. But it is not clear that adequate draining of the core can occur through passive means alone. Moreover, thermal spectrum MSR lack the inherent negative temperature reactivity feedback of LWRs. This can result in positive reactivity feedback and power instabilities.

On balance, there is no compelling evidence at this point to support the claim that MSR will be safer than LWRs. On the contrary, there are many characteristics that would present additional and potentially severe safety challenges.

Sustainability

One significant potential advantage of liquid-fuel MSR compared to LWRs is improved sustainability. As discussed in chapter 3, some MSR reportedly would use uranium more efficiently, generate less long-lived nuclear waste, and even use existing nuclear waste from LWRs as fuel. Below, the sustainability benefits of several MSR concepts compared to LWRs are considered, and found to be modest at best. But to realize even a modest benefit, MSR fuels would require some form of reprocessing, with its attendant proliferation and security risks.

MOLTEN SALT BREEDER REACTOR

One of the original motivations for developing MSR in the 1960s was their potential to operate as thermal breeder reactors when fueled with U-233 and thorium (see chapter 2). While a fast reactor is required to breed plutonium-239 from U-238, in theory a thermal reactor can be used to breed the fissile isotope U-233. This is because in a thermal spectrum, fission of U-233 releases more neutrons than U-235 or plutonium-239. These extra neutrons could be used to bombard fuel containing the isotope thorium-232 to produce more U-233 than is needed to maintain the chain reaction. This is a process similar to the breeding of plutonium-239 in fast reactors from bombarding U-238 with neutrons.

However, there is a catch that makes U-233 breeding difficult or even impossible in conventional, solid-fueled reactors: the production of the isotope protactinium-233 (see Box 8). That problem can be circumvented only by rapidly reprocessing the fuel as the reactor is operating to remove the protactinium-233. Such “on-line” reprocessing can only be carried out at a reactor with liquid fuel. Thus, the molten salt breeder reactor (MSBR) concept was born. In principle, the MSBR is more sustainable than the LWR.

URANIUM UTILIZATION EFFICIENCY

A recent Oak Ridge National Laboratory review analyzed the fuel cycle for a 1000 MWe MSBR starting up with a core containing thorium and 19.75 percent-enriched HALEU (Gehin and Powers 2016). The use of HALEU for the startup core would be necessary since U-233 does not occur naturally, and must be produced by irradiating Th-232. The study assumed a scenario in which after startup, the reactor would be able to breed enough fissile U-233 to be self-sufficient, and would require only additions of thorium—no additional HALEU would be needed. The reactor would also be operated for 30 years, over which time the graphite moderator would have to be replaced seven times. In this idealized situation, the quantity of natural uranium and thorium needed per GWe-yr, averaged over a 30-year lifetime, would be 36 metric tons, compared to the 180 metric tons of natural uranium per GWe-yr that an LWR needs: a five-fold improvement.²⁴

LONG-LIVED WASTE GENERATION

TRU generation is lower in a reactor fueled with thorium and U-233 than in a reactor that contains U-238, because more neutron absorptions are required to generate plutonium-239 and heavier transuranic isotopes. According to the Oak Ridge

study, for the idealized MSBR, the quantity of TRU waste that would be generated per GWe-yr that would need to be disposed of in a repository at the end of the 30-year operating lifetime, assuming an annual discharge to waste of 4.7 percent of the steady-state TRU inventory of the core, would be only around 3.3 kilograms (kg), compared to about 230 kg/GWe-yr for an LWR (Gehin and Powers 2016).²⁵ This is about a 70-fold decrease, which, going forward, would not qualify as a significant reduction according to the National Academy of Sciences standard, but would meet the DOE's long-lived waste reduction goal of at least a factor of 10 (see chapter 3).

TECHNICAL CHALLENGES

However, other recent analyses of the MSBR have highlighted many technical challenges with the concept that would make it highly impractical for use as a commercial reactor. For example, the idealized MSBR discussed above, which starts up with a HALEU core but only requires refueling with thorium afterward, is unphysical. Studies have found that an MSBR that starts up with HALEU instead of HEU or U-233—an essential nonproliferation requirement—would need to operate at least 20 years to reach a steady-state (Betzler, Powers, and Worrall 2017; Zou et al. 2018). During

BOX 8.

Protactinium and the Thorium Fuel Cycle

For a nuclear reactor to be an effective breeder, neutrons must be used very efficiently to convert fertile material, such as U-238, to fissile material, such as plutonium. One needs to minimize losses of neutrons by absorption in non-fertile materials. In a thorium-based breeder reactor, where the goal is to breed the fissile isotope U-233 by neutron bombardment of thorium-232, the generation of the intermediate product protactinium-233 is a problem that must be addressed if efficient breeding is to occur.

After U-238 absorbs a neutron and is transmuted to U-239, it decays with a half-life of about 24 minutes to neptunium-239, which then decays with a half-life of about 2.4 days to plutonium-239. Because neither of these intermediate isotopes is around for a long time, it is unlikely that either will absorb another neutron before eventually decaying to fissile plutonium-239. And if the neptunium-239 does absorb a neutron, it will decay into plutonium-240, which is also useful in nuclear fuel.

However, the situation is more challenging if one wants to breed U-233 by irradiating thorium-232 with neutrons. In that case, an intermediate isotope is created, protactinium-233, with a half-life of 27 days. If protactinium-233 absorbs a neutron before decaying to U-233, it will become U-234,

which is not useful for nuclear fuel. Because protactinium-233 has such a long half-life, there is a high likelihood if it stays in the reactor that it will absorb a neutron and thus will not decay to U-233, degrading the reactor's capability to breed new fuel.

Solving the protactinium problem and developing a workable thermal breeder reactor using thorium fuel was one of the chief motivations for the original MSR project at Oak Ridge National Laboratory. To effectively breed U-233, protactinium-233 would be separated from the liquid fuel with an on-line reprocessing system as soon as it was generated and stored outside of the reactor until it had decayed to U-233.

However, this separation process poses a proliferation danger because it produces pure U-233, which is a weapon-usable isotope comparable in risk to plutonium-239. To address this concern, researchers developed the concept of a "denatured" MSR. In such reactors, protactinium-233 is not removed from the salt, and LEU is added, which dilutes the U-233 produced. The resulting uranium in the reactor is also low-enriched and less attractive for nuclear weapons than separated U-233. However, these reactors would not be breeders.

that transition period, HALEU would have to be fed into the reactor until enough U-233 was bred for self-sustaining operation, which would greatly reduce the sustainability performance of the system.

One study identified a potential transition scenario in which a HALEU-fueled MSBR would be operated for the period of time necessary to produce enough U-233 for a new startup core and subsequent refueling over the remaining lifetime of the reactor (Yu et al. 2019). The reactor would have to be fed HALEU for three years in order to produce enough U-233 for the replacement core (by separating and storing Pa-233, as discussed in Box 8, p. 95). Each additional year of the initial operation with HALEU would produce enough U-233 for twenty years of operation without HALEU. An MSBR with a 30-year lifetime would therefore have to operate for about 4.3 years with additions of HALEU. The natural resource efficiency (uranium plus thorium) for this system would be about 2 percent, or about three times that of an LWR—and 60 percent lower than for the unphysical MSBR described earlier.

Similarly, about 140 kg/GWe-yr of plutonium and other TRU would be generated during the time that HALEU is fed into the system, or about 600 kg in total. This additional TRU would contribute to the overall TRU generation of the steady-state MSBR, estimated above at about 75 kg, for an average TRU generation of around 30 kg/GWe-yr. This is about eight times less than the average TRU generation rate for an LWR—again, not nearly as impressive as the reduction for the idealized MSBR with no transition period to steady-state operation, and not significant according to the factor-of-10 standard.

Another obstacle is the need for an efficient on-line reprocessing system to at least produce sufficient U-233 to regenerate the fuel that was consumed, and thus allow self-sustaining operation (see Box 8, p. 95). To break even, the reactor fuel would have to be reprocessed at least once every three days to remove the accumulating protactinium-233 (Gehin and Powers 2016). It is not clear that such a chemical processing system would be feasible. The process would entail a complex series of steps, most of which have not been demonstrated beyond laboratory scale. Researchers have observed that “some of the separation processes are considered too difficult to be implemented” (Mathieu et al. 2006).

In spite of these problems, variations of the molten salt breeder reactor concept are being pursued by the startup companies Fluor Energy in the United States and Seaborg Technologies in Denmark. As with the original design of this type, both of these reactors would require continuous reprocessing. Seaborg Technologies bills its reactor as a “waste-burner” that would use TRU obtained by reprocessing LWR

spent fuel, although a 2015 technical white paper on its website suggests that fresh LEU should be used, at least for its pilot plant (Seaborg Technologies n.d.) (It is not clear whether this company is aware of the studies showing the problems with using LEU as a startup fuel (Betzler, Powers, and Worrall 2017; Zou et al. 2018).

THE DENATURED MSR

Because of its need for on-line reprocessing, the molten salt breeder reactor would pose a high proliferation risk—something the US government recognized as far back as the 1970s. The concerns are two-fold. First, the need to produce and store large quantities of weapon-usable U-233, poses security risks, and second, the presence of the on-line reprocessing system makes the reactor much more difficult to safeguard than an LWR.

In response to these concerns, the DOE developed a concept for a 1000 MWe “denatured” MSR that (1) would not have an on-line chemical separations capability, and (2) would dilute the fissile uranium isotopes (U-233 and U-235) with U-238 to a concentration comparable to that of LEU. (Because U-233 has a smaller critical mass than U-235, U-233 must be diluted to less than about 12 percent in a mixture with U-238 to render the material impractical for use in nuclear weapons, compared to 20 percent for U-235.)

Since the denatured MSR would not have on-line reprocessing, it would never be able to breed its own fuel, even at a steady state, and thus would require a constant supply of fresh fissile fuel, such as 19.75 percent-enriched HALEU, to compensate for the buildup of neutron-absorbing fission products. (As with all MSRs, gaseous fission products and some metallic fission products would still have to be removed from the denatured MSR’s liquid fuel, and plant operators would have to manage the resulting radioactive wastes.)

URANIUM UTILIZATION EFFICIENCY

While these additional features make the denatured MSR less proliferation-prone than the MSBR, they also make it less resource-efficient. Although the denatured MSR would be able to achieve a higher fuel burnup than an LWR, as discussed in chapter 3 the use of higher-assay LEU requires a greater quantity of natural uranium to produce each batch of fuel. The only way to increase natural uranium utilization in a once-through fuel cycle is by extending burnup by increasing the internal conversion of fertile to fissile fuel (in this case, from thorium-232 to U-233), which reduces the need for fissile material additions. But without on-line reprocessing, the reactor cannot produce new fissile material quickly enough to greatly reduce the demand for additional fresh fuel, and thus needs a steady supply of HALEU.

Consequently, the denatured MSR's annual resource requirements (for natural uranium and thorium) would be about 2.5 times less than the uranium requirements for an LWR, a more modest reduction than the factor of five achieved by the idealized MSBR (Gehin and Powers 2016), but comparable to a more realistic MSBR system that takes the transition period into account.

LONG-LIVED WASTE GENERATION

Because of the need to add U-238 to denature the U-235 and U-233, the denatured MSR produces more plutonium and other long-lived TRU. As a result, instead of the 75-fold reduction in TRU waste that the MSBR would achieve relative to an LWR, the denatured MSR would have only about a four-fold reduction (Gehin and Powers 2016). This is below the DOE's factor-of-10 significance criterion.

THORCON

The ThorCon reactor is similar in concept to the denatured MSR. ThorCon's intent is to develop a small modular reactor that can be rapidly commercialized (the company calls it the "do-able" MSR). To do so, the design utilizes currently available materials and technologies to the extent possible—constraints that further limit its sustainability performance.

URANIUM UTILIZATION EFFICIENCY

"On a once through basis, ThorCon is not that uranium efficient," according to an archived page that used to be on the company's website (ThorCon 2017). Information provided by the company's principal engineer, Jack Devanney, illustrates why (Jack Devanney, email message to the author, January 4, 2018.). A 1000 MWe ThorCon plant, consisting of four 250 MWe reactor modules, would require 55 metric tons of thorium and 9.44 metric tons of 19.75 percent HALEU to start, and an addition of 7.2 kg of HALEU per day during eight years of operation. This corresponds to an average requirement of about 3.8 metric tons of HALEU and 6.9 metric tons of thorium per year. To produce this much HALEU would require about 160 metric tons of natural uranium per year, for a total natural resource requirement of about 167 metric tons per year, compared to about 180 metric tons per year of natural uranium for a similarly sized LWR. This is a reduction of less than 10 percent.

According to ThorCon, one of the reasons why the reactor is less uranium-efficient than the denatured MSR is that each 250 MWe ThorCon module is smaller and uses a proportionately smaller amount of thorium, thereby producing less U-233 (Jack Devanney, email message to the author, January 4, 2018). This means that more HALEU must be added each year to keep the reactor operating. Also, the

fuel needs to be replaced every eight years, as opposed to 30 years for the denatured MSR.

The only way the ThorCon design could utilize uranium more efficiently overall is if the spent fuel salt, discharged after eight years, were sent to a reprocessing plant to recover the remaining fissile materials for reuse as fresh fuel. Of course, this would be inconsistent with the chief rationale for the denatured MSR and undo its nonproliferation benefits.

LONG-LIVED WASTE GENERATION

What about waste generation? Like the denatured MSR, ThorCon would generate long-lived TRU in its spent fuel at a lower rate than LWRs because of its use of thorium. However, without reprocessing and recycling its spent fuel salt, a ThorCon plant would not generate less long-lived nuclear waste than an LWR.

A 1000 MWe ThorCon plant would generate an average of about 27 metric tons of spent fuel salt per year over its eight-year core lifetime, compared to an average of about 22 metric tons of spent fuel generated per year by an LWR over a 60-year lifetime (ThorCon 2018). The volume of this waste would be approximately nine cubic meters, the same as the nine cubic meters of LWR spent fuel discharged annually.

ThorCon salt waste would contain approximately 11 percent by weight uranium, with a U-235 enrichment of 5.2 percent and a U-233 enrichment of 3.4 percent. This is a higher concentration of fissile uranium isotopes than in LWR spent fuel, although this mixture of uranium isotopes would be considered "low-enriched" (Forsberg et al. 1998). It would also contain about 1.5 percent by weight plutonium and other TRU—a slightly higher concentration than for typical LWR spent fuel.

Because of the large quantity of unused fissile material in the spent salt waste, ThorCon proposes a sequence of reprocessing steps, using gaseous separation methods, to remove the leftover uranium from the spent fuel salt and convert the remainder into an ash form. The recovered uranium would be reused. ThorCon claims that after these processes, the volume of the salt waste would be reduced by more than 80 percent. But the weight fraction of plutonium and other TRU in this concentrated waste would be greater than the weight fraction in LWR spent fuel. Combined with the fact that the ThorCon waste would contain less cesium-137, which provides the self-protecting radiation barrier in LWR spent fuel (see chapter 4), the ThorCon waste would be more attractive from a proliferation perspective.

Like other thorium-fueled reactors, the ThorCon reactor would produce a smaller quantity of TRU overall than an LWR. A 1000 MWe plant would generate about 50 kg of TRU per year, compared to 220 kg for an LWR—about the same reduction factor as the denatured MSR. Neither reactor would meet the DOE factor-of-10 significance criterion.

Overall, it does not appear that the ThorCon reactor would be significantly more sustainable than the LWR.

INTEGRAL MOLTEN SALT REACTOR (TERRESTRIAL ENERGY)

The IMSR being developed by the company Terrestrial Energy shares some features of the previous MSR designs. However, it differs from them in that it uses only LEU with a U-235 content below 5 percent (that is, no thorium or HALEU). Each IMSR400 module would have a power rating of approximately 440 MWth and would generate “up to 195 MWe” of electrical power (Terrestrial Energy 2020).²⁶ The reactor’s first core would contain an LEU-based fuel salt enriched to less than 2 percent U-235 and would use makeup fuel salt enriched to 4.95 percent U-235 (Choe et al. 2018). Each module would be used for seven years—the lifetime of the graphite moderator—and then would be discarded and swapped out for a new one. In total, the reactor would operate for eight cycles, or 56 years.

The IMSR would use LEU fuel with an enrichment comparable to LWRs, but the maximum achievable burnup would be only about one-third as high as typical LWR fuel. Thus, the reactor would not be expected to be more sustainable than an LWR, as is borne out by the below discussion.

URANIUM UTILIZATION EFFICIENCY

The IMSR, operating on a seven-year once-through cycle, would require an average of 277 metric tons of natural uranium per GWe-yr—around 50 percent more than an LWR (Choe et al. 2018). The designers have proposed an alternative fuel cycle in which the spent fuel salt is used to start up new reactor modules, which would increase its uranium utilization. However, assuming this is even technically possible, it would reduce the average natural uranium requirement to 194 metric tons per GWe-yr, still greater than the LWR requirement (Choe et al. 2018).

LONG-LIVED WASTE GENERATION

Similarly, Terrestrial Energy predicts that the IMSR will generate about 175 kg of plutonium per GWe-yr, compared to about 200 kg per GWe-yr for an LWR (Choe et al. 2018). Using its proposed alternative fuel cycle, this would decrease to about 105 kg of plutonium per GWe-yr. In either case, the reduction would not be considered significant according to the DOE’s factor-of-10 criterion.

TRANSATOMIC POWER

Transatomic Power was founded in 2011 to develop a uranium-fueled, moderated MSR that it originally called the Waste-Annihilating MSR (see Box 9, p. 99). As the name indicated,

the company’s major selling point was the reactor’s sustainability: Transatomic claimed its reactor could “annihilate” nuclear waste by running entirely on spent nuclear fuel (Zanolli 2015). In addition, it said that the reactor could also use LEU fuel 75 times more efficiently than LWRs (Temple 2017). The Transatomic design differed from ThorCon in that it would not use thorium, and it differed from the Terrestrial Energy IMSR in that it would utilize an extensive on-line reprocessing system to remove fission products to achieve the very high burnup it projected.

However, an independent review of the concept conducted at the Massachusetts Institute of Technology found that there were serious errors in the company’s calculations and that the reactor could not maintain a chain reaction by using the TRU from spent fuel as its feedstock (Temple 2017). Further analysis by Oak Ridge National Laboratory, in conjunction with Transatomic Power, also showed that the reactor would use uranium far less efficiently than originally claimed (Robertson et al. 2017). Transatomic Power then abandoned the notion that its reactor could run on spent fuel and significantly scaled back its claims about increased uranium utilization.

URANIUM UTILIZATION EFFICIENCY

The 2017 Oak Ridge National Laboratory review of the Transatomic Power reactor found that if the reactor used 5 percent-enriched uranium feed, it could only operate for 29 years, reaching a fuel burnup of less than 10 percent. Over its lifetime, it would achieve a uranium utilization rate of about 1 percent (Robertson et al. 2017). This is only marginally better than the 0.6 percent efficiency of an LWR.

LONG-LIVED WASTE GENERATION

After 29 years of operation, the 520 MWe Transatomic Power reactor operating on 5 percent-enriched LEU would contain nearly 3 metric tons of plutonium (Robertson et al. 2017).²⁷ This corresponds to an average plutonium production rate per GWe of about 200 kg of plutonium per year, which is about the same as for current-generation LWRs. Therefore, there is no advantage with respect to this metric.²⁸ The Transatomic Power reactor would actually produce plutonium at a greater rate than an LWR for the first 15 years of operation. The annual rate of plutonium production in the first five years would be about 670 kg per GWe—more than three times the LWR’s rate of plutonium production.²⁹

In addition to the generation of high-level waste, one must also consider the generation of low-level and TRU waste. The on-site reprocessing system would generate much more of these wastes than an LWR. The Oak Ridge analysis concluded that “the [low-level waste] associated with the

continued on p. 100

Transatomic Power: A Cautionary Tale

Transatomic Power famously burst onto the scene in 2011 with a TED (Technology, Entertainment, Design) talk by its founders, Mark Massie and Leslie Dewan, introducing the Waste-Annihilating MSR. The company heavily promoted the idea that its reactor could convert spent fuel from LWRs into clean energy (Zanolli 2015), using this aspect as a selling point in radio advertisements sponsored by the Nuclear Energy Institute. The company also initially claimed that the reactor could use 1.8 percent-enriched LEU as its fuel, could run for 100 years, and would be able to “generate up to 75 times more electricity per ton of mined uranium than a light-water reactor” (Temple 2017).

There are technical reasons why these claims merited skepticism from the outset. As an MSR with a moderated neutron spectrum, it was unclear how the design could effectively fission the TRU in spent fuel. As discussed in chapter 3, in thermal reactors TRU are more likely to absorb neutrons than fission, reducing the reactivity of the fuel and making the neutron chain reaction harder to maintain. Moreover, in the Transatomic Power reactor, TRU would build up in the fuel because there was no plan to extract them from the salt by on-line reprocessing. However, because the average neutron speed in this reactor would be somewhat faster than in an LWR, the designers claimed that there were sufficient higher-energy neutrons to effectively fission TRU.

The extremely high uranium utilization claim was also suspect. It corresponds to a fuel burnup of 96 percent, compared to around 5 percent for LWRs. As discussed here, this was a far higher utilization than was predicted for other MSRs, including the molten salt breeder reactor, which is counterintuitive. A 96 percent burnup means that 96 percent of all the uranium loaded into the reactor, including the 65 metric tons of uranium fuel in the initial core, would be ultimately fissioned and converted to energy. Burning up this fraction of the initial core alone would generate enough power to run the reactor for over a century. But the 500 kg of fresh fuel added each year would also have to be burned almost completely. Thus the claim made little sense from the beginning. Nevertheless, Transatomic received very favorable media attention and attracted the interest of venture capitalists.

In late 2016 all references to nuclear waste as a fuel source for its reactor were removed from the company’s website. It then stated its goal was no longer to “reduce existing stockpiles of spent nuclear fuel” but “to reduce nuclear waste

production by significantly increasing fuel burnup” (TAP 2016a). Accordingly, the company then focused on a version of its reactor that would use 5 percent-enriched LEU and would operate for less than 30 years. And its claims for the resource efficiency of this design became far more modest. Instead of an improvement of a factor of 75, Transatomic said that the design allowed for only “more than twice” the fuel utilization of LWRs (TAP 2016a).

This drastic scaling down of Transatomic’s claims for its design were made after an external review conducted at the Massachusetts Institute of Technology revealed serious errors, as revealed publicly in a Technology Review article in early 2017 (Temple 2017).

As Transatomic put it in the final iteration of its technical white paper:

This version of the [Transatomic Power] design white paper incorporates multiple revisions based on further research performed over the past year. . . . [T]his work has revealed new understandings about the system . . . [, and] . . . we have realized that our initial analyses of spent nuclear fuel core loadings were centered around inaccurate assumptions about reactor behavior that . . . had to be corrected (TAP 2016b).

Subsequently, the company was referred to as the “Theranos of nuclear power” (Baron 2017)—after the infamous company that marketed a finger-prick blood test that did not really work—and it began keeping a much lower profile. The Nuclear Energy Institute stopped profiling Transatomic Power in its radio advertisements. And the company’s chief executive officer, Leslie Dewan, posted a video on YouTube in which she explained how to learn from failure.

Transatomic Power continued operating after these missteps became public, but it was never able to right itself. Although the design itself was no less viable than those being developed by other MSR startups, the company’s backtracking likely led to a loss of confidence by investors, and it finally shut down in September 2018. Transatomic’s experience should serve as a warning to other nuclear reactor vendors not to promise more than they can reasonably deliver. These companies should also strive to make public as many details as possible about their systems, to ensure that their claims can be subject to rigorous peer review.

continued from p. 98

continuous or batch processing of a large amount of fuel” is a disadvantage (Robertson et al. 2017).

Thus, it does not appear that the Transatomic Power reactor could in fact have provided a large enough benefit in sustainability to justify the cost of development and the additional safety and proliferation risks—a realization that was likely a factor in the company’s demise.

FAST-NEUTRON MSRS

Given the many challenges and limitations of thermal-spectrum MSRs, some experts have concluded that the concept is not worth pursuing. Instead, they believe that fast MSRs, which would not require a graphite neutron moderator, are more promising (Mathieu et al. 2006). As with solid-fuel fast reactors, a fast spectrum MSR is better than a thermal MSR for converting fertile material to new fuel and has a greater tolerance for neutron-absorbing fission products and TRU. For these reasons, fission products can be removed at a much lower rate from fast MSRs than from thermal MSRs without affecting the breeding potential, thus requiring a less intensive on-line reprocessing system.

Fast MSR designs have even been proposed that may be able to operate for more than 20 years without the need for on-line fuel reprocessing or fresh fuel additions (Mathieu et al. 2009). Others might be able to operate longer by functioning as breed-and-burn reactors with no reprocessing requiring enriched uranium at startup but needing only depleted uranium as feed (Hombourger et al. 2015). Such reactors, if feasible, would be able to utilize uranium more efficiently than thermal MSRs such as the ThorCon or Transatomic Power reactors. And, unlike Transatomic Power and other slower-neutron MSRs, fast MSRs could in principle be fueled with the TRU extracted from spent LWR fuel by reprocessing—that is, they could burn some components of nuclear waste.

After years of review, researchers in France decided to limit that country’s MSR research to developing a thorium-fueled, fluoride-based fast reactor called the Molten Salt Fast Reactor (MSFR). In the United States, chloride-based fast MSRs are being pursued by a TerraPower-Southern Company consortium (as a complement to TerraPower’s solid-fueled fast reactors) and the startup Elysium Industries. The DOE also has made a bet on fast-spectrum MSRs, selecting Southern Company for an award of up to \$90.3 million to build a test reactor called the Molten Chloride Reactor Experiment.

However, this shift in focus has been criticized by supporters of thermal MSRs, who argue that fast neutron variants will require more development work and hence are

less “doable” (Zwartsenberg 2016). As they note, fast-spectrum MSRs have not been demonstrated, even at an experimental scale. And structural materials that can survive bombardment by lots of fast neutrons in a corrosive molten salt environment have not even been developed yet.

In addition, recent studies have highlighted the technical limitations of fast MSRs and cast further doubt on their ability to operate more sustainably than LWRs in practice. First, a fast MSR requires a large amount of fissile material—at least five metric tons of U-233 or more than 12 metric tons of TRU at startup. (Brovchenko et al. 2019). Therefore, in order to establish and support an expanding fleet of fast MSRs, the reactors must breed and separate large quantities of fissile fuels to sustain their own operations and produce fuel for new reactors. And even though the required rate of on-line reprocessing would be lower than for thermal MSRs, it is still significant: around 0.2 percent of the fuel would have to be reprocessed daily, with tons of weapon-usable materials separated and recycled each year.

Second, the necessary reprocessing technologies have not been demonstrated or even fully defined in some cases. For instance, chloride salt-based fast MSRs may require a pyroprocessing technology similar to that being used for the EBR-II spent fuel. But as discussed in chapter 5, that process has had major problems and would need significant performance improvements to be useful for an MSR. And small-scale experiments on processing the molten fluoride salts that would be used in the European MSFR have found only a “low” extraction efficiency for uranium (Rodrigues, Durán-Klie, and Delpech 2015).

Third, for chloride-based fast MSRs such as the TerraPower design, isotopically pure chlorine-37 most likely would be needed for advantageous nuclear characteristics (Napier 2020). This would require enrichment of natural chlorine, 75 percent of which is the undesirable isotope chlorine-35. Chlorine-37 is not currently commercially available in bulk quantities (Napier 2020), but is only sold in milligram-sized quantities of sodium chloride. It is unclear how long it would take and how much it would cost to establish a bulk supply. One study stated that chemical methods for chlorine enrichment would be “unattractive for cost reasons,” and speculated that laser enrichment (which is not currently available) could produce this material at a reasonable cost, but does not provide an estimate (Hombourger et al. 2015). In an admittedly unscientific survey, this author obtained a quote in February 2018 from a US chlorine-37 supplier of \$46 per milligram at 98 percent enrichment. Given that billions of times this quantity would be required for a fast MSR, one can see why the current price is “unattractive.”

A 2019 analysis of the MSFR provides some insight on its potential for improved sustainability, although it uses some artificial simplifying assumptions (Brovchenko et al. 2019). The 1500 MWe reactor can be started up with a core consisting of 30.6 metric tons of natural thorium and 12.8 metric tons of TRU obtained from reprocessed LWR spent fuel. After 60 years of operation, the core inventory of TRU decreases to around 800 kg.

URANIUM UTILIZATION EFFICIENCY

To produce enough TRU to fuel the core of the 1500 MWe MSFR, 58 1000 MWe LWRs would have to operate for one year, requiring about 10,440 metric tons of natural uranium. (The initial amount of natural thorium required for the MSFR is relatively small and can be omitted from these calculations.) After 60 years of MSFR operation, the combined LWR-MSFR system would generate a total of 148 GWe-years (58 GWe-years of LWR operation plus 90 GWe-years of MSR operation). Assuming the MSFR could operate for 60 years without the need for additional fissile fuel, the system would require an average of about 70 metric tons of natural uranium per GWe-year—a factor of 2.5 less than an LWR. This is not a significant improvement and is not much better than what the Transatomic Power moderated MSR could have achieved.

Chapter 8 discusses the potential for fast MSRs to operate in a self-sustaining breed-and-burn mode, which could increase their uranium utilization efficiency without the need for online reprocessing (other than removal of gaseous and insoluble fission products).

LONG-LIVED WASTE GENERATION

After 60 years of operation of the MSFR, the amount of remaining TRU—about 800 kg—can be compared to the amount that an LWR of the same capacity would have generated had it operated instead. A 1.5 GWe LWR of LWRs operating for 60 years would generate about 19.8 metric tons of TRU. When this is added to the quantity of TRU that would have been used for the first core of the MSFR, the result is 32.6 metric tons. Thus, the TRU reduction factor would be $32.6/0.8$, or about 40—which would meet the DOE's factor-of-10 criterion, but not the more stringent factor-of-100 standard.

In practice, the reduction factor would be smaller. The 2019 MSFR analysis does not consider the process losses resulting from on-line fuel reprocessing (Brovchenko et al. 2019). If a 1 percent process loss per cycle is assumed (an optimistic assumption), then more than 1 metric ton of unrecovered TRU would end up in the waste stream over 60 years of operation. Taking this loss into account would lower the TRU reduction factor to only 11.

ADDITIONAL MSR WASTE ISSUES

Another downside of MSRs relative to LWRs is that molten salt radioactive waste has physical and chemical characteristics that make it far more difficult to manage and dispose of than the solid spent fuel generated by LWRs. This has been shown by real-world experience—the failure to effectively dispose of the waste from the MSRE at Oak Ridge National Laboratory, which shut down nearly 50 years ago. The legacy waste from this experiment has proven to be one of the most persistent and technically challenging cleanup problems in the DOE complex.

After the reactor was shut down in 1969, nearly five metric tons of spent fuel salt, along with a similar quantity of slightly contaminated flush salt, was dumped into storage tanks and allowed to cool and solidify. There were no plans for management of this material beyond long-term interim storage. Despite attempts to control the generation of gases due to radiolysis (radioactive decomposition of the salt), high concentrations of highly hazardous uranium hexafluoride and fluorine gas were detected more than 25 years later. The spent fuel salt presented severe chemical and radiological hazards, and it also had the potential to cause an accidental criticality, largely due to the presence of U-233.

The DOE decided in 1998 to remove all of the waste from the drain tanks. However, that proved to be too difficult, partly because some of the fuel had solidified and clogged the piping. Ultimately, the DOE was able to remove most of the U-233 through a chemical separation process, but this left the residual radioactive salts behind. Further work to remove them was suspended. The radiation doses in the waste tanks are still around 1,000 rem/hour, far too high to allow human access. The salt also continues to generate fluorine gas, which must be pumped out of the tanks and treated every six months in a system that has been experiencing failures recently after 20 years of service and must be replaced. Given the potential hazards of further cleanup operations, the DOE is now seriously considering entombing the remains of the reactor and waste tanks in concrete, creating a permanent repository in situ (Huotari 2017; McMillan 2019).

The viability of any new nuclear reactor concept depends critically on whether it can be safely decommissioned and its nuclear waste effectively managed. In light of that, Oak Ridge National Laboratory's failed experience with trying to clean up the mess left behind by the MSRE contradicts the often-heard notion that this demonstration project was successful.

Commercial MSR operators would have to manage and dispose of hundreds of tons of waste salt, far more than the five metric tons generated by the relatively small MSRE. In

addition to spent fuel salt, operators would have to manage at least six other distinct waste streams, including captured noble gases and their decay products (Riley et al. 2019). MSR proposals cannot be regarded as credible unless they include feasible plans for managing and disposing of these wastes.

Proliferation/Terrorism Risk

As discussed above, to maximize fuel burnup and sustainability, liquid-fueled MSRs must have chemical processing systems to remove fission products from the fuel on a frequent basis. In other words, such MSRs must have their own reprocessing plants and are thus not simply nuclear reactors, but are reactors with on-site fuel cycle facilities. Reactors with co-located reprocessing facilities would pose proliferation risks. A recent preliminary proliferation study of the European MSFR concluded that while it would be “impossible” to divert nuclear materials directly from the reactor, “it would be possible to do so by misusing the salt cleaning [i.e., reprocessing] facility” (Allibert et al. 2020).

Thermal-spectrum MSRs such as the Molten Salt Breeder Reactor or the Transatomic Power reactor would require multiple reprocessing cycles per year. In contrast, fast-spectrum MSR fuel may not have to be reprocessed as frequently because it could potentially tolerate higher fission product levels. However, this benefit would be offset to some extent by the larger quantities of fissionable materials that fast MSRs require.

Some MSR concepts, such as ThorCon, would not have on-site reprocessing, but the company assumes that spent fuel salt would be sent off-site for reprocessing to recover unused fuel. In addition, LWR reprocessing facilities would be needed if MSRs such as the European MSFR are to burn TRU from LWR spent fuel.

The potential of some MSRs to achieve a high conversion ratio raises the possibility that some designs may be able to operate as breed-and-burn reactors, as discussed in chapter 8. However, because MSRs present challenges for material accountancy, these reactors would have greater proliferation risks than solid-fueled breed-and-burn reactors operating on a once-through cycle.

While the specifics of the required chemical fuel treatment processes at MSRs vary from one design to another, they would all present significant challenges for nonproliferation and nuclear security.

MATERIAL ACCOUNTANCY AND SAFEGUARDS

MSRs are on-line fueled reactors like the pebble-bed HTGRs discussed in chapter 6, but they are even more difficult to safeguard. MSR fuel is not contained in discrete and countable items such as HTGR pebbles or LWR fuel

assemblies. With regard to material accountancy, an MSR is more like a bulk-handling fuel cycle facility than a conventional nuclear reactor. As discussed in chapter 4, bulk-handling facilities are especially challenging to safeguard. But in contrast to bulk-handling facilities where material is fed into the process in batches, the fuel in MSRs continuously circulates through the reactor and—for some designs—also through a co-located reprocessing facility. This makes timely detection of diversions of fissile material even more difficult because material inputs and outputs are harder to define and measure.

Keeping track of the fissile material inventory in an MSR would also be a challenge because the material would be distributed throughout the system “in more locations in more forms,” requiring a “substantial increase in instrumentation complexity” (Qualls and Holcomb 2019). Fissile materials could be transported with the off-gas and deposited onto various reactor surfaces (Qualls and Holcomb 2019). And as discussed in chapter 4, the pyroprocessing technologies needed for fuel treatment are not very efficient, resulting in the discarding of significant quantities of fissile materials in hard-to-measure waste streams.

These difficulties are compounded by the fact that the material flow rate through the on-line reprocessing system at a single MSR would be considerably larger than at a centralized reprocessing plant designed to handle the spent fuel from many reactors. This is problematic for material accountancy because the measurement uncertainty is proportional to the facility throughput.

The Transatomic MSR design is a good illustration of those challenges. (Transatomic Power was considerably more transparent about the details of its design than many other reactor startups, and it has now made all its intellectual property available to other researchers.) The 520 MWe reactor, fueled by LEU, would operate for 29 years without interruption. After startup, the plutonium content of the molten salt in the core rises steadily to a peak of about 4 metric tons after 20 years of operation, after which it slowly decreases to about 3 metric tons at shutdown (Robertson et al. 2017). During the first part of its 29-year operating cycle, an intermediate-energy neutron spectrum promotes conversion of U-238 to plutonium. In the second part of the cycle, a thermal neutron spectrum promotes fission of the plutonium that has built up in the core.

To achieve a 29-year lifetime, the MSR core would have to be processed to remove neutron-absorbing lanthanide fission products approximately every 50 days—which means that the entire core would flow through the reprocessing plant more than seven times a year. The corresponding flow rate of plutonium through the plant would be as high as 80 kg per day, or more than 29 metric tons per year (Robertson et al.

2017). To put this in context, the reprocessing facility for a single 520 MWe Transatomic Power MSR would need to process more than three times as much plutonium per year as a large industrial plant such as La Hague in France—which is capable of handling the spent fuel discharge of 40 1,000 MWe LWRs annually. This is a challenge because the greater the amount of plutonium flowing through the plant per year, the more difficult it is to accurately account for it.

Fast-spectrum MSRs would require lower reprocessing rates than thermal reactors, as discussed above. However, they would still have very high fissile material flow rates because their cores are large. For example, the initial TRU core of the 1500 MWe MSFR would contain about 11 metric tons of plutonium, which would cycle through the reprocessing plant every 450 days, corresponding to an initial plutonium throughput of about nine metric tons per year (Brovchenko et al. 2019).

In 2014, the IAEA pointed out that more stringent nuclear material accountancy methods will be needed for liquid-fueled reactors, and the instrumentation by and large remains to be developed (IAEA 2014). However, such development work has been slow, and the statement remains true today. A 2018 review by the national laboratories concluded that “the IAEA and its international safeguards system do not have the policies, concepts, approaches, or technologies needed for applying safeguards to MSR designs” (Kovacic et al. 2018).

SAFEGUARDING PYROPROCESSING AT MSRS

The online reprocessing plants at MSRs could not use aqueous technologies such as PUREX but would use pyroprocessing technologies, which are more challenging to safeguard (see chapter 4). The challenges of safeguarding pyroprocessing plants would be even greater at MSRs. These would require co-located plants employing different types of pyrochemical separation processes, depending on system requirements. Because these plants would operate on a continuous basis, it would be even harder than at a batch-loaded pyroprocessing plant to keep track of the material flowing through them and separating into various product and waste streams.

There is little public information about the details of the fuel processing systems that MSRs would require, including factors relevant to material accountancy, such as the anticipated uncertainties in calculating the quantities of radionuclides in the core, the efficiency of separating fuel and waste streams, and the amount of fuel expected to be contained in the process’s residual holdup. In fact, because the separation processes have not been demonstrated on a commercial scale or in some cases even tested, it is not known whether they would be possible, much less feasible on a commercial scale.

In 2014, a review article pointed out that fundamental data for the extraction processes are lacking, especially for

the separation of actinides from lanthanide fission products (Serp et al. 2014). Subsequently, experimental work on separating uranium and neodymium (a surrogate for TRU) from molten fluoride salt found only “low” extraction efficiencies, calling into question the proposed processing approach (Rodrigues, Durán-Klie, and Delpech 2015).

In any event, it is highly unlikely that an effective safeguards approach based on material accountancy could be developed for the pyroprocessing systems at MSRs. Even with a process loss rate of 0.1 percent per year, which would be remarkably low, more than one SQ of plutonium would be discharged into waste streams every year. Accurate accounting for this material would be difficult and costly.

As is the case for pyroprocessing plants more generally, the safeguards approach for MSRs would likely be even less reliant on material accountancy than conventional fuel cycle facility safeguards, and would instead depend more on complementary measures such as containment and surveillance and process monitoring. But, as discussed in chapter 4, such measures cannot entirely compensate for a lack of accurate material accountancy, and if the IAEA ultimately accepts them as substitutes, the risk of diversions could increase. Moreover, the lack of timely material accountancy also presents security concerns, as it could prevent plant operators from quickly determining whether a terrorist claim of theft were true. The safeguards and security risks of the very high fissile material production and processing rates of MSRs may well prove unmanageable and ultimately disqualify these reactors from widespread deployment.

RADIOACTIVE XENON RELEASES AND COMPREHENSIVE TEST BAN TREATY VERIFICATION

MSRs could also create problems for the nonproliferation regime by emitting noble gas fission products that could be mistaken for the radiological signatures of underground nuclear weapon tests (Lyman 2019). These emissions could interfere with the International Monitoring System set up to help verify the Comprehensive Test Ban Treaty (after it enters into force) by detecting clandestine nuclear weapon tests. The radioactive xenon emissions from medical radioisotope production plants are already causing background noise today that is reducing the effectiveness of the system. As a result, in 2015 the United Nations asked producers of commercial medical isotopes to reduce, and, if possible, eliminate their releases of radioactive xenon. Scientists have identified a target level of xenon emissions that nuclear facilities should keep below in order to avoid unacceptable interference with the International Monitoring System. However, meeting this goal has proven to be difficult (Jubin, Paviet, and Bresee

2016). Expensive and cumbersome off-gas capture and delay systems would be required.

A single 440 MWth Terrestrial Energy IMSR would generate a thousand times more xenon-133 per day than a radioisotope production facility, and 10 million times as much as the target level (Lyman 2019). Unless MSR designs incorporate the required off-gas systems to achieve emissions reductions to this level—which will likely be costly and difficult—deployment of only a handful of these reactors over the world could significantly interfere with Comprehensive Test Ban Treaty verification (Lyman 2019).

Readiness for Commercial Demonstration and Deployment

There is a wide range of opinion on how soon any MSR design could be commercialized and brought to market. On the low end, the Canadian company Terrestrial Energy predicts that it could bring its IMSR to commercial markets in the 2020s (Terrestrial Energy n.d.). Terrestrial Energy bases its optimistic outlook on what it calls its reliance on “proven and demonstrated” MSR technology (Terrestrial Energy n.d.). In early 2017, it announced that it planned to submit an application for design certification or a combined license to the Nuclear Regulatory Commission in 2019, although it ultimately failed to meet that timetable. The company is continuing to engage in pre-application interactions with the agency, partially supported by around \$500,000 in DOE funding.

Other MSR startups, such as ThorCon and Elysium, also claim they will have a commercial product available by the mid-2020s. Kairos Power, which is developing a solid-fuel MSR referred to as a fluoride high-temperature reactor, received funding from the DOE in 2020 to build a reduced-scale test reactor, and is aiming for a commercial demonstration by 2030.

A different view was expressed by the DOE’s 2017 advanced demonstration and test reactor study, as discussed in chapter 1. The study judged that even the fluoride high-temperature reactor, which would use a solid TRISO fuel, is a low-maturity technology that requires “significant research, development, and demonstration” before it could be commercialized (Petti et al. 2017). The DOE assessed that this reactor would first require an engineering demonstration that would take 10 to 15 years to begin and cost \$2 billion to \$4 billion.

The DOE concluded that the overall technological readiness of liquid-fueled MSRs was comparable to that of the solid-fueled fluoride high-temperature reactor, despite the fact that liquid fuels are less mature than TRISO fuel. However, the report pointed out that liquid-fueled reactors

would be harder to license and to safeguard against diversion of nuclear materials. The report assessed that these reactors would not be commercially available before 2045 to 2050—two decades after Terrestrial Energy’s aggressive deployment date—and did not provide cost estimates for the additional development and demonstration work needed to reach that point (Petti et al. 2017). The main difference between the development timelines of the DOE and the MSR startup companies is that the DOE believes that additional performance demonstrations are required before commercial demonstration projects can move forward.

The French safety agency Institut de Radioprotection et de Sûreté Nucléaire (IRSN) has been even more pessimistic than the DOE, stating that “it seems hard to imagine any [molten salt] reactor being built before the end of the [21st] century” (IRSN 2015).

Whether one believes the projections of the vendors or those of the DOE and IRSN, commercialization of any MSR will most likely require a substantial investment of time and resources—moving the goal out of reach of most private ventures without a significant infusion of government support. As with the other NLWR designs, one must ask what advantages such reactors would bring, and whether these benefits would justify the development costs.

PRIOR DEMONSTRATIONS AND THEIR RELEVANCE

Whose time scale is more realistic—MSR vendors who claim their designs could be commercialized within a decade, or the DOE, who maintains that it would take several decades? As for other advanced reactor concepts, the timeline largely depends on whether a design has been successfully demonstrated at both engineering and commercial scales—demonstrations that incorporate the major operational and safety features of the proposed commercial plant to the extent possible.

Several startups believe the experience gained from the MSRE in the 1960s was sufficient to allow them to leapfrog over additional engineering or performance demonstration steps and proceed to commercial demonstration. However, the MSRE design of the 1960s is significantly different from the MSR designs being considered today. Moreover, the data sets collected during the MSRE have major gaps and are of limited use and relevance for the development and licensing of commercial MSRs.

APPLICABILITY OF THE MSRE TO THE TERRESTRIAL ENERGY IMSR

The example of the Terrestrial Energy IMSR400, which is currently in pre-application review by the Nuclear Regulatory Commission, is instructive. The company says that the IMSR

design was “based heavily” on the MSRE, but it has not clearly shown how the MSRE experience was sufficient or even applicable. In response to a recent pre-licensing submittal document to the Nuclear Regulatory Commission, the agency pointed out that although Terrestrial Energy noted that its design builds off the MSRE work done at Oak Ridge National Laboratory, the submittal apparently did not indicate exactly where the lab’s work was considered; it also did not identify any differences between the IMSR and the MSRE (NRC 2020b).

The MSRE was not similar enough to the IMSR to serve as a demonstration that the IMSR could run safely and reliably. The MSRE design, capacity, fuel salt composition, and operational parameters, as well as its system for managing the radioactive off-gases and other wastes, were substantially different from those of the IMSR.

Fuel Characteristics and Off-Gas Generation

Each IMSR module would have a power rating of 440 MWth, about 60 times greater than the MSRE. The fuel would contain LEU with a U-235 content below 5 percent, instead of the HEU and U-233 used in the MSRE. As a result, the fuel salt would have to have a much higher concentration of total uranium to achieve criticality than the MSRE fuel had, which would affect numerous reactor and fuel properties. Also, while Terrestrial Energy is vague in public about the fuel salt chemical composition the IMSR would use, it has said that the fuel does not include either lithium or beryllium salts—two of the MSRE fuel constituents (LeBlanc 2021). While Terrestrial Energy is wise to avoid the use of beryllium—an extremely toxic metal—it cannot rely on the MSRE’s experience to support use of a different type of fuel.

In any event, the MSRE did not shed much light on how to qualify MSR fuels more generally. As Oak Ridge scientists recently pointed out,

MSRE experience provides limited guidance as to what fuel salt properties would be necessary to measure at future MSRs. MSRE did not operate its fuel salt to high burnup and did not need to measure changes in thermochemical properties (Holcomb, Poore, and Flanagan 2020).

The maximum burnup attained by the U-235-containing fuel salt in the MSRE was about 13,250 MWd/MTHM. The IMSR fuel burnup, although relatively low, would be somewhat higher than that (Choe et al. 2018).

The MSRE also did not provide sufficient data to validate models of the behavior of isotopes such as xenon-135 that have a critical impact on reactor operation. According to a recent survey, “no models of MSR xenon behavior that can

be reasonably said to be validated exist” (Price, Chvala, and Taylor 2019). And unlike the MSRE and most MSR designs, the current IMSR concept would not have a means of systematically removing the radioactive xenon that is generated in the fuel, but would allow the xenon gas to accumulate until it is naturally released.³⁰ Accurately modeling the impacts of xenon behavior under such an approach will need considerable experimental validation.

Safety Systems

Terrestrial Energy claims that its reactor would not need any active systems to maintain safety (IAEA 2016). However, the IMSR would not use the conventional MSR approach of employing a freeze plug that would melt in the event of a loss of cooling accident, enabling the fuel to drain from the reactor into special dump tanks. Instead, it would rely on passive natural convection cooling of the core (LeBlanc 2016)—a significant difference from the MSRE, which used a freeze plug. As is also the case for sodium-cooled fast reactors and high-temperature gas-cooled reactors, as discussed in previous chapters, such passive cooling systems have not been demonstrated for commercially sized units.

And as discussed above, the DOE has stated that “the need for active safety systems cannot be ruled out” (Petti et al. 2017). Thus regulators may decide that the current IMSR design would have to be significantly modified by adding active safety systems before it could be licensed.

Given the absence of an MSR performance demonstration that is sufficiently representative of the IMSR, regulators should require a prototype for performance demonstration before the IMSR could be licensed and commercially deployed—as the Nuclear Regulatory Commission expected in the 1990s for sodium-cooled fast reactors and high-temperature gas-cooled reactors.

Reactor Materials

The MSRE also experienced many problems with reactor materials, including corrosion of metallic reactor structures caused by high-temperature molten salts, and irradiation-induced swelling of the MSRE’s graphite that would necessitate its periodic replacement. For example, MSRE data suggested that graphite would need to be replaced every three to five years (Busby et al. 2019). However, the data are not very useful because MSRE’s graphite was a different grade than contemporary graphite. Moreover, there are flaws with much of the material property data that were collected, such as poor-quality photographic images (McFarlane et al. 2019). Data for corrosion and neutron irradiation of structural materials were sparse (McFarlane et al. 2019). The incomplete data are of limited value for critical activities such as

developing materials that can perform adequately in MSR environments.

Terrestrial Energy's approach to the materials issues observed at the MSRE, which is similar to ThorCon's, would require significant technical review. Instead of pursuing development of structural materials that would be more robust under the punishing conditions in MSRs, the IMSR designers would limit the service lifetime of its modular reactor units to seven years—although, given the MSRE experience, it is not even clear that the graphite could be safely used for that long. To accomplish this, Terrestrial Energy has proposed the use of “Core-units” that would consist of the fuel, pumps, heat exchangers, and graphite moderator, all sealed within a reactor vessel (Terrestrial Energy 2020). These Core-units would be swapped out for new ones after reaching their seven-year design lifetime. To achieve a reactor lifetime of 56 years, comparable to today's LWRs, each Core-unit in an IMSR would have to be replaced eight times. And since it would take about six IMSR modules to replace the capacity of a single large LWR, 48 Core-unit replacements would be needed to generate the equivalent amount of electricity over a normal plant lifetime, whereas the LWR would require none.

Terrestrial Energy's plan is for each Core-unit to be shipped to the reactor site sealed; they would not need to be opened by the reactor operator for any reason. Fresh fuel would be pumped continuously into the core and spent fuel would be stored within each vessel. It is unclear how the gaseous and noble metal fission products that would be removed from the fuel, including the cesium-137 waste discussed earlier, would be managed and disposed of, as Terrestrial Energy does not plan to use the MSRE off-gas recovery and treatment system. (In any event, that system would not provide an adequate basis for MSRs today because it experienced significant difficulties (McFarlane et al. 2019). The off-gas system was also cumbersome and is not likely to be feasible for commercial MSRs if they are to achieve significant reductions in noble gas emissions (Lyman 2019).)

After seven years, the vessels containing the Core-units would be replaced, and the old ones would remain in on-site storage for years until they could be moved to an undetermined final disposal site. The reactor vessels would not be small and would be difficult to handle: each is 7 meters tall, with a diameter of 3.6 meters, and a weight of 170 metric tons. Although the company portrays this as a simple approach, management of these vessels could become a safety and logistical nightmare for a utility. Moreover, given the uncertainties in materials performance, it seems premature to design a

sealed reactor vessel that workers could not access for inspection and maintenance.

In addition to fundamental questions of reactor design, there are many other issues that would need to be resolved before the IMSR could be commercially deployed. For example, Terrestrial Energy does not say where and how the fuel—the composition of which is a commercial secret—would be manufactured. The resources needed to finance, locate, design, license, and build a plant to manufacture this unique fuel are likely to be very large. Finally, as is the case with all the MSR startups, Terrestrial Energy does not appear to have a well-formulated plan for management and disposal of the IMSR's spent fuel. This critical aspect of MSR operation is routinely given short shrift.

APPLICABILITY OF THE MSRE TO OTHER MSR DESIGNS

Other MSRs under development, such as the TerraPower fast MSR, have designs that are even more different from the MSRE than the IMSR. Notably, these would use a chloride-based molten salt instead of the fluoride-based salt used in the MSRE. Since chloride salts have “little or no irradiation performance data and are generally more corrosive,” (McDuffie 2017), such a reactor will require significant research and development compared to fluoride-based MSRs.

IMPACT OF SAFEGUARDS ON MSR DEPLOYMENT

Even if MSRs were technically mature, the lack of an international safeguards approach could prove to be a stumbling block to near-term deployment. For instance, Terrestrial Energy is seeking to site its first commercial unit in Canada, a non-nuclear weapon state. But even though the baseline IMSR design would not have on-site reprocessing, monitoring the continuous uranium input and accounting for spent fuel in sealed vessels that cannot be inspected or assayed would also prove difficult for safeguards. The project cannot proceed far until the IAEA develops the necessary techniques and protocols for applying material accountancy and other verification measures at MSRs.

This study finds that, given the unresolved safety and security issues, aggressive timelines to commercially deploy MSRs within a decade are infeasible and would compromise safety. The DOE's more conservative view—that MSRs would not be ready until at least the 2040s—is much more realistic and would allow more thorough resolution of the many safety, security, and environmental issues raised by reactors with liquid fuels.

Breed-and-Burn Reactors

The Rationale for Breed-and-Burn Reactors

SUSTAINABILITY AND REPROCESSING

As discussed in chapter 3, one of the main arguments for establishing a closed fuel cycle—reprocessing spent fuel and recycling recovered materials usable as new nuclear fuel in reactors—is that the once-through cycle is “unsustainable.” Advocates of reprocessing say that because light-water reactors (LWRs) convert only about 0.5 percent of natural uranium to energy, they do not use uranium efficiently and will rapidly deplete the world’s uranium resources. They also assert that LWRs generate large quantities of nuclear wastes that pose long-term radiological and security risks.

Many non-light-water reactor (NLWR) vendors claim that by incorporating reprocessing into their fuel cycles their designs can burn existing nuclear wastes and/or use uranium more efficiently. The waste argument in particular has recently gotten more traction as most countries with nuclear power, including the United States, have failed to make progress in building geologic repositories for disposal of spent nuclear fuel. However, as shown in chapter 3, the potential for reprocessing and recycle systems to increase sustainability over reasonable timescales is greatly exaggerated. In practice, it is highly unlikely that closed fuel cycles could significantly reduce the need for uranium or effectively solve the nuclear waste problem. In any event, the world is not in danger of running out of uranium any time soon, as discussed in chapter 1, and uranium remains so cheap that there is little economic incentive to reduce its use. Most importantly, as discussed in chapter 4, reprocessing and recycling raise serious proliferation and nuclear terrorism concerns. The marginal

sustainability benefits of fuel cycles with reprocessing and recycling are far outweighed by their proliferation and nuclear terrorism risks.

Nevertheless, some maintain that the LWR operating on a once-through cycle is a wasteful dead-end for nuclear power. To be sure, though some economists may disagree, there is some merit to the principle that natural resources should be used more efficiently, even those that are cheap and plentiful today. And, as discussed in chapter 3, developing more uranium-efficient reactors is one approach to reduce the environmental impacts of uranium mining.

So it is worth considering whether technological advances can improve the sustainability of nuclear power without the need to reprocess and recycle spent fuel. This line of thinking has led to a renewed push for development of once-through “breed-and-burn” reactors. Successful development of such reactors could be advantageous for security because they could undercut the most compelling motivations for the closed fuel cycle. However, although some progress has been made, many technical and economic challenges remain for the development of safe and effective breed-and-burn reactors.

CAN NUCLEAR POWER BE MADE MORE URANIUM-EFFICIENT WITHOUT REPROCESSING AND RECYCLING?

More than 99 percent of natural uranium is the isotope uranium-238 (U-238), which is much less likely to undergo fission than U-235 or plutonium-239 when struck by a neutron. To use natural uranium more efficiently than today’s LWRs in a once-through cycle, a reactor system would need to be able to convert a larger fraction of the U-238 component of natural

uranium to plutonium-239 and then fission the plutonium to produce energy, before a given load of fuel reaches the end of its useful life.

There are two primary constraints on the amount of U-238 that LWRs can convert to plutonium-239 and fission to release energy on a once-through basis. The first is the burnup of current generation nuclear fuels, which is related to the length of time that they can be used in a reactor, and is limited by both physical and neutronic changes that occur during irradiation. The second is the inability of LWRs to use depleted uranium waste from enrichment directly as fuel. To greatly increase natural uranium utilization in the once-through cycle, both these constraints must be addressed.

INCREASE URANIUM UTILIZATION BY INCREASING PHYSICAL LIMITS ON FUEL BURNUP

As uranium-based fuels are irradiated, some of the U-238 is converted to plutonium-239, and some of that plutonium is fissioned to produce energy. However, not enough plutonium is produced to make up for the decrease in U-235 due to fission and the accumulation of fission products that capture neutrons but do not fission. Also, certain isotopes of plutonium and other transuranic (TRU) elements are more likely to capture a neutron rather than undergo fission, depending on the neutron energy. As a result, over time the fuel no longer has enough fissile material to maintain a chain reaction. Fuel burnups are also limited because irradiation changes the physical structure and material properties of the fuel pellets and cladding. The metal cladding also corrodes from contact with cooling water and fission products. Eventually, the fuel deteriorates to the extent that it can no longer be used safely and must be discharged as waste.

If LWR spent fuel is reprocessed, then the residual plutonium that has not been fissioned (about 1 weight-percent) could be separated and used to manufacture new fuel for LWRs or NLWRs. However, more than 90 percent of the spent fuel is U-238, which is not useful as a fuel unless it is converted to plutonium. For this reason, simply reprocessing spent fuel and recycling plutonium does little to increase uranium efficiency.

Current-generation LWR fuel stays in the reactor for five to six years at most, and average fuel burnups are around 50,000 megawatts-days per metric ton (MWd/MT). At that burnup, about five percent of the initial uranium in the fuel either undergoes fission directly (the U-235 and a small fraction of U-238) or is converted to plutonium and fissioned (the U-238). To increase LWR fuel burnups, new materials, production processes, and fuel loading patterns would be needed for fuel to be safely used for longer periods of time. As discussed in chapters 5, 6, and 7, some types of NLWR

fuels, in principle, can reach much higher burnups than current-generation LWR fuels.

INCREASE URANIUM UTILIZATION BY INCREASING FUEL SHUFFLING

The core of an LWR is subdivided into a few (usually three) batches. During a refueling outage, which occurs every 18 months to two years, one of the batches is removed from the core and replaced with fresh fuel. The other two batches are then rearranged, or “shuffled.” This is done both to improve safety and to increase uranium utilization by more effectively using the neutron flux in the core for power production.

Uranium utilization could be increased further by subdividing the core into a greater number of batches and increasing the number of times and/or frequency that fuel is shuffled (Xu 2003). (This is the principle behind current-generation CANDU reactors, which are designed to be refueled while operating. They are about 30 percent more uranium efficient than LWRs, but their natural uranium fuel cannot achieve high burnups.) Reactors with the capability to load and shuffle fuel frequently or even continuously have the potential for more efficient uranium use than reactors with conventional batch refueling. For instance, as discussed in chapter 6, the on-line-refueled pebble-bed high-temperature gas-cooled reactor (HTGR) is about 33 percent more uranium-efficient than the prismatic-block HTGR, although it is not necessarily more uranium-efficient than the LWR (Bays and Piet 2010).

The ultimate on-line-fueled reactor is the MSR, which allows for continuous refueling, adjustment of fuel composition, and reprocessing to extract neutron-absorbing fission products. The main advantage of MSRs is that in theory they can reach very high burnups without running up against the physical limits of solid fuels and cladding materials. However, as discussed in chapter 7, the use of liquid fuel introduces a host of other problems, such as the continuous release of fission product gases from the molten fuel. And MSRs designed to maximize uranium utilization must employ on-line reprocessing, and therefore are not once-through systems, although some concepts would not require chemical separations but only physical removal of gaseous and insoluble fission products.

INCREASE URANIUM UTILIZATION BY IRRADIATING DEPLETED URANIUM

The second constraint on uranium utilization is the inability to use the large fraction of mined uranium that does not end up as fuel but is discarded as enrichment process waste. Even if 100 percent of LWR reactor fuel could be fissioned and

converted to energy, that would amount to only about 10 percent of the amount of natural uranium mined to produce the fuel.

The remaining depleted uranium, primarily U-238, is contained in the enrichment waste, i.e., the uranium “tails.” In order to increase the uranium utilization above 10 percent, a reactor system would need to be able to use the depleted uranium tails as fuel. The challenge is that depleted uranium has little reactivity of its own, so it cannot be used as a reactor fuel without being supplied with extra neutrons from another source (for instance, the driver fuel in fast breeder reactors, as discussed in chapters 2 and 3). And burnup of the depleted uranium fuel would also be constrained by the same factors that limit other fuels (e.g., irradiation damage and low reactivity).

Because of this uranium tails problem, simply increasing fuel burnup does not necessarily increase the efficiency of uranium utilization. Increasing fuel burnup typically requires using a higher level of uranium enrichment to increase the amount of U-235 in the fuel. But this presents a catch-22 situation since the additional enrichment needed would result in more depleted uranium waste. Therefore, for the once-through cycle to use uranium more efficiently, a reactor system would need to (1) extend fuel burnup without the need to increase the level of uranium enrichment, (2) use U-238 as fuel in steady-state operation without the need for additional input of U-235 or other fissile isotopes, or (3) both. Reactors that could increase the internal conversion of U-238 on a once-through basis, reducing the need for external supply of fissile isotopes, have been studied for decades. Some reactors would have mixed fast and thermal neutron spectra within the same core. There have even been attempts in the past to develop so-called spectral shift LWRs, in which the initial neutron spectrum would be slightly faster to optimize conversion of U-238 to plutonium-239 and then later would be slowed down and made more thermal to optimize plutonium fission.

Such higher-conversion reactors would still require regular refueling to load additional fissile material into the core and thus would only meet criterion (1) above. However, these reactors would only have modestly greater uranium utilization efficiency—perhaps 50 percent at most. Larger increases would require a reactor that could achieve both (1) and (2). In principle, such a “breed-and-burn” reactor would use uranium much more efficiently without reprocessing and recycling, would require less uranium enrichment, and would reduce the amount of spent fuel generated. A reactor that could achieve these goals on a once-through basis would negate the main rationale for reprocessing and recycling spent fuel.

Breed-and-Burn Reactors

The defining property of a breed-and-burn reactor is the capability of achieving a once-through steady-state mode of operation that can utilize U-238 as fuel without requiring additional inputs of fissile material such as U-235. A number of breed-and-burn reactor designs have been proposed in past decades. Most recently, the concept saw a revival with TerraPower’s pursuit of a traveling wave reactor. Most of these designs are liquid metal-cooled fast reactors, primarily because they can supply extra neutrons to breed enough plutonium to enable self-sustaining operation. However, fast MSRs, as well as fast gas-cooled reactors such as General Atomics’ EM2, have been studied for their potential to operate in breed-and-burn mode.

TECHNICAL CHALLENGES

Unfortunately, breed-and-burn reactors do not appear to be any simpler to develop than any other type of NLWR. In fact, these reactors are even more technically challenging, because the requirements for high burnup, once-through fueling, and the use of U-238 impose additional constraints on design and operation.

Safety is a paramount concern. Breed-and-burn reactors challenge the safety limits of conventional reactors, almost by definition. Achieving breed-and-burn operation requires fuel burnups far exceeding the current experience base and the known limits of existing materials. In addition, as discussed in chapter 2, breed-and-burn reactors must have very low leakage of neutrons from the core, to maximize the number of neutrons available for converting U-238 to plutonium. They also must generate a large fraction of their power from fission of the plutonium that is produced. For sodium-cooled breed-and-burn fast reactors such as TerraPower, these requirements lead to large, positive void reactivity coefficients, making them even less stable than conventional fast reactors and necessitating the addition of novel safety systems which have not yet been developed (Qvist and Greenspan 2012).

Positive void coefficients and their associated instabilities also proved to be a significant issue in LWR spectral shift reactors because of the higher plutonium inventories in the reactor. Designs that avoided this problem only increased uranium use efficiency by 20 percent—an insignificant increase—compared to ordinary LWRs (Martin et al. 1991).

Some observations related to the sustainability, safety, and security issues of breed-and-burn reactors are presented below. Breed-and-burn reactors must be safe and reliable if they are to be viable nuclear power options for the future. The payoff could be large if the sustainability of the once-through cycle can be significantly improved, undercutting

the primary rationale for reprocessing and recycling. However, there are many significant technical obstacles in their development path, and it is unclear whether these can be overcome. Developing reactors that can achieve the multiple goals of enhanced safety, sustainability, and good economics will be a major challenge.

Breed-and-Burn Reactor Designs

TRAVELING WAVE REACTOR

The breed-and-burn reactor design that has attracted the most attention is the traveling wave reactor, first proposed by H-bomb pioneer Edward Teller. This type of reactor resembles a candle: a “wick” of fissile material, such as high-assay low-enriched uranium (HALEU), is used to initiate a chain reaction at one end of a column of U-238. A wave front would then travel slowly along the column. As the wave front advances through the core, successive layers of U-238 would first be converted to plutonium (“bred”) and then fissioned (“burned”). The hope was that, once begun this process would be stable and sustainable, enabling the reactor to operate for decades without the need for refueling.

A number of variants of this concept have been proposed in the last few decades. However, the only one that has been seriously pursued for commercialization is the TerraPower traveling wave. The TerraPower design is a sodium-cooled, metal-fueled fast reactor modeled after the Experimental Breeder Reactor II (EBR-II) (see chapter 5).

While the original concept for the TerraPower reactor was a candle-like traveling wave reactor, after the company began to advance the design beyond preliminary feasibility studies, it discovered that the traveling wave approach was not likely to work (Hejzlar et al. 2013). TerraPower then adopted a more conventional “standing wave” design in which the breed-and-burn wave remains stationary but the fuel in the reactor vessel is periodically shuffled. (Despite this major change, TerraPower continued to refer to its design as a traveling wave reactor, noting that the wave still travels in the reference frame of the fuel as it is shuffled, although not with respect to a stationary observer.)

The reactor core has a burn zone consisting of HALEU driver fuel assemblies (up to 19.75 percent U-235) and a breed zone in which depleted uranium feed assemblies are loaded. The driver and feed assemblies are shuffled every 18 months to reduce peaks in the power distribution, which enables them to be irradiated more efficiently. As plutonium is bred in the feed assemblies, they are moved to central core positions to replace spent driver fuel assemblies. In order to maintain a lifetime core without the need to refuel the reactor vessel, a

zone within the vessel is provided for internal storage of fresh feed assemblies and spent fuel assemblies (Kim and Taiwo 2010).

Other institutions working on breed-and-burn designs, such as the Korea Advanced Institute of Science and Technology, are intent on pursuing approaches that do not require shuffling, due to concerns about the potential for mishaps resulting from fuel movements (Yonghee Kim, professor in the Department of Nuclear and Quantum Engineering, Korea Advanced Institute of Science and Technology, conversation with the author, May 6, 2016). However, TerraPower’s experience suggests that these concepts may prove infeasible in practice.

TECHNICAL OBSTACLE: DEVELOPMENT OF HIGHER BURNUP FUELS

Another technical challenge for breed-and-burn reactors is the development of fuels that can reach very high burnups. Depleted uranium feed assemblies must remain in the reactor long enough to build up a sufficient concentration of plutonium and then function as driver assemblies. To sustain breed-and-burn operation in a TerraPower-type fast reactor, researchers have calculated that a batch of fuel must attain an average burnup of around 20 percent (that is, fission of 20 percent of the initial heavy metal nuclei). Moreover, because the power distribution is not uniform in these reactors, peak fuel burnup must reach 30 percent in order for the average burnup of a batch of fuel to reach 20 percent (Greenspan 2016).

These burnups exceed the historical irradiation experience for this type of fuel. This is true both for the uranium-zirconium alloy fuel and for the steel cladding material known as HT9, which has only been demonstrated experimentally to maintain its integrity to about 10 percent burnup in fast test reactors such as the EBR-II. The damage to the cladding from bombardment by fast neutrons in a breed-and-burn reactor would be more than twice as great as the peak damage in EBR-II fuel (Hejzlar et al. 2013). Also, irradiating fuel to such high burnups results in high fission gas pressure that puts strain on the cladding and could cause it to fail. The standard solution to this problem is to provide large plenums (empty spaces) in the fuel rods into which fission gas can expand, but this is apparently not sufficient to lower pressure on the cladding enough to achieve the burnups needed for breed-and-burn.

Accordingly, a major focus of TerraPower had been developing fuel and cladding materials that can sustain high enough burnups to make breed-and-burn operation possible. This entailed painstaking and time-consuming experimental work, such as irradiating more than 1000 cladding samples

in the BOR-60 fast reactor in Russia. But to address the fission gas problem, TerraPower also had to adopt vented fuel pins that continuously release fission products directly into the coolant. These fission products must be extracted from the coolant and either stored or released, creating additional problems for safety, environmental protection, and radioactive waste management.

Another way to address the burnup limitation is to develop reactor cores that would allow fuel to be shuffled in three dimensions instead of only two (along the vertical axis as well as the two horizontal axes). This would allow for an even more finely tuned matching of fuel and local neutron flux. Researchers have developed modified breed-and-burn fast reactors that would use fuel rods subdivided into segments (Hou, Qvist, and Greenspan 2015). These segmented rods could be shuffled vertically as well as horizontally. This approach would reduce the minimum required peak burnup by 30 percent and nearly double the uranium utilization. However, the segmented fuel rods could not have gas plenums and, therefore, like the TerraPower fuel, would have to allow for venting of fission product gases into the coolant. Moreover, even this approach would require fuel capable of withstanding radiation damage nearly twice the current demonstrated level. There are no simple fixes for these complex problems.

These technical challenges may have proved too formidable for TerraPower. The company, which was founded in 2006, initially intended to build a 600 megawatt-electric (MWe) demonstration plant in China as early as 2022, and in 2017 signed a joint venture with the China National Nuclear Corporation to complete the design and commercialize the technology. However, the project was cancelled after the Trump administration imposed additional restrictions on the export of nuclear technology to China in 2018, and TerraPower then said it hoped to build a demonstration traveling-wave reactor in the United States. However, it was clear that a near-term demonstration reactor would not have been able to operate in breed-and-burn mode, in part because the company had not yet solved the fuel burnup problem. Idaho National Laboratory wrote at the time that “TerraPower is proceeding with the first prototype while acknowledging that achieving their ultimate design goals in terms of high burnup fuels with high fluence cladding will require additional testing beyond the first prototype” (INL 2018). In other words, the demonstration reactor would have been a conventional, metal-fueled fast reactor requiring periodic refueling.

More recently, TerraPower’s goal of building a conventional demonstration fast reactor became closer to realization when the DOE chose the Terrapower-GE Hitachi 345 MWe Sodium design for the Advanced Reactor Demonstration

Program (ARDP). Around the same time, the company said that it was pausing development of the traveling wave reactor (Freebairn 2020). As discussed in chapter 5, the Sodium will be less uranium-efficient than an LWR. Demonstration of actual breed-and-burn operation will likely not occur unless fuels capable of attaining ultra-high burnups are fully developed and qualified—potentially decades away.

PEBBLE-BED BREED-AND-BURN REACTORS

Another design with three-dimensional fuel shuffling that has been proposed is a metal fuel pebble-bed fast reactor, which could be cooled either by gas or by liquid metal. (The graphite used as a matrix material for pebbles in thermal gas-cooled reactors would have to be replaced with a less moderating material in a fast reactor.) According to a University of California, Berkeley, study, this approach could smooth out burnup variations throughout the core and would use uranium more efficiently than designs with only two-dimensional shuffling (Greenspan 2016). Unfortunately, the metal-fueled pebble bed reactor concept ran into various problems, and the Berkeley group halted work on it (Greenspan 2016).

MOLTEN SALT BREED-AND-BURN REACTORS

As discussed in chapter 7, some types of MSR may be able to operate as breed-and-burn reactors. Recent work indicates that breed-and-burn is only possible with chloride salt-based fast MSRs (using chlorine enriched to nearly 100 percent chlorine-37) with a uranium-plutonium fuel cycle (Martin et al. 2017). But such reactors must be very large in physical size, because neutrons are not effectively blocked by chlorine and thus are more likely to escape from the core. In addition, the density of plutonium and other TRU in the fuel is limited by the properties of the salt, which also leads to larger core volumes. For example, one reactor of this type analyzed would require 432 metric tons of uranium contained in about 860 metric tons of salt and have a power rating of 33,000 MWe—more than 20 times the power rating of the largest commercial LWR, the 1600 MWe Evolutionary Power Reactor (M.V. Martin, undergraduate student in the Department of Bio-engineering, University of California, Berkeley, conversation with the author, January 30, 2018). Such a design is clearly not suitable for deployment as a small modular reactor, and it is unlikely any utility would want to buy such a large reactor.

Breed-and-burn MSRs with lower power ratings are theoretically possible. However, for a reactor with a reasonable power rating, such as 1000 MWe, uranium fuel fed into the reactor would have to remain in the reactor for hundreds of years before it reached a burnup that could sustain breed-and-burn operation (M.V. Martin, undergraduate student in

the Department of Bioengineering, University of California, Berkeley, conversation with the author, January 30, 2018). Researchers have not yet calculated how long such reactors would have to run on enriched uranium before they could achieve breed-and-burn operation and be fueled on depleted uranium alone, but it is likely to be a long time for reactors of a feasible size.

It has also not been determined how high a burnup the fuel could actually achieve, which would determine the reactor's lifetime. The reactor cannot achieve breed-and-burn operation unless the fuel can achieve a sufficiently high burnup. Unlike the traveling wave reactor, the MSR's fuel burnup would not be limited by degradation of the solid fuel structure. But without reprocessing, fission products would accumulate in the salt, and eventually the concentration of waste products would get so high that the reactor could no longer operate. The impact of this issue on the feasibility of the concept is unclear (Hombourger et al. 2015).

In addition, manufacturers would have to be found that could supply the huge quantities of enriched chlorine needed to make the fuel for these reactors. As discussed in chapter 7, this material is not currently commercially available and could be a huge expense.

Therefore, although breed-and-burn operation may be possible in theory for certain molten salt fast reactors, such reactors appear completely infeasible for commercial deployment.

Sustainability

URANIUM UTILIZATION COMPARED ACROSS REACTOR TYPES

One of the main motivations for developing breed-and-burn reactors is to utilize uranium more efficiently without reprocessing. How well would the current concepts do in this respect compared to current once-through reactors? And how would they compare to the uranium utilization of a closed fuel cycle with fast breeder reactors?

URANIUM EFFICIENCY OF ONCE-THROUGH BREED-AND-BURN REACTORS VS. LWRs

As noted above, the uranium utilization efficiency of current-generation LWRs is only about 0.6 percent, and, although the burnup of LWR fuel could be increased by using higher-enriched fuel containing more U-235, this alone would not improve uranium utilization. To achieve that goal, more U-238 in the fuel must also be converted to plutonium and fissioned, a process constrained by reactor design.

In contrast, breed-and-burn reactors could utilize uranium more efficiently by being able to use depleted

uranium as fuel, which would increase the average fuel burnup without increasing the enrichment required over the reactor's lifetime. In a 2010 report, DOE researchers estimated the uranium utilization efficiency for a number of different once-through reactors, including early traveling wave reactor concepts and the TerraPower reactor (Kim and Taiwo 2010). The calculated utilization efficiencies ranged from 0.9 percent to 29.4 percent, and the highest values assumed average fuel burnups of nearly 30 percent.

One limitation of this DOE study is that it only evaluated the reactor's uranium utilization at equilibrium (steady-state breed-and-burn operation), when it would be able to operate using only depleted uranium feed. In practice, however, the reactor would have to operate for many decades and use a considerable amount of HALEU fuel before it could reach equilibrium. This means the reactor must be initially loaded with a larger inventory of HALEU than would be required simply to make it critical. Taking the transition to equilibrium operation into account, current breed-and-burn reactor concepts would be only a few times more uranium-efficient than LWRs over a realistic time scale.

For example, TerraPower estimates that 32 metric tons of natural uranium on average would be needed per 1000 MWe-yr for the TWR, assuming a 60-year lifetime (Gilleland, Petroski, and Weaver 2016). (Since the reactor would not be refueled, its lifetime would be limited by the amount of fuel that could be initially loaded into the reactor vessel, in addition to other factors such as the lifetime of reactor structures.) The amount of uranium fissioned per year would be around 0.9 metric ton (less than the amount for an LWR because fast reactors are slightly more thermally efficient than LWRs). The average uranium utilization would be 0.9 divided by 32, or around 3 percent, roughly five times that of the LWR. As mentioned earlier, the uranium utilization could be doubled with a core design with segmented fuel assemblies that would allow three-dimensional shuffling, to about 6 percent. But as is the case for a fast breeder reactor system with reprocessing and recycling, the overall uranium utilization could not be greatly increased during the lifetime of a single reactor because a large stockpile of depleted uranium would have been generated to produce the initial core (although not as large as that needed for a plutonium breeder).

TerraPower gets around this problem by suggesting the possibility of transferring entire cores to a second generation of reactors once the first-generation reactors reach the end of their lifetimes, and so on (Gilleland, Petroski, and Weaver 2016). But this assumption requires the same leap of faith as the assumption that a closed-cycle fast breeder reactor system will be completely rebuilt as many times as necessary

to use up all the uranium that was originally mined to start up the system. If we accept the assumption, however, the uranium utilization efficiency will increase over time and eventually becomes equivalent to the maximum burnup of the depleted uranium fuel, which would be around 20 percent for the metal fast reactor fuel that the TWR would use.

URANIUM UTILIZATION EFFICIENCY OF ONCE-THROUGH BREED-AND-BURN REACTORS VS. CLOSED-CYCLE FAST BREEDER REACTORS

How does the uranium utilization efficiency of once-through breed-and-burn reactors compare to closed-cycle fast breeder reactors? At first glance, not so well. Breed-and-burn does not come close to the nearly 100 percent uranium utilization that can theoretically be achieved by plutonium-fueled fast breeder reactors in a closed fuel cycle. However, as discussed in chapter 3, the 100 percent figure is misleading because it can only be achieved if the fast breeder reactor system and its associated fuel cycle operate for tens of thousands of years, with zero loss of material to reprocessing waste. Over a 60-year lifetime, the uranium utilization of the fast breeder at equilibrium would only be about 0.3 percent—that is, 10 times worse than the TerraPower reactor.

Therefore, the uranium utilization efficiency of breed-and-burn systems appears to be as good as, or even better than, fast breeder reactors with reprocessing and recycle, at least over a time scale of centuries.

Although breed-and-burn reactor fuel would have a higher burnup than LWR fuel, the spent fuel would still contain a substantial inventory of uranium, plutonium, and other TRU. One study estimates the total TRU discarded to waste per GW-year is about two times that of LWRs (Diego di Sanzo 2014). Some breed-and-burn concepts would further increase uranium utilization by incorporating what is referred to euphemistically as “reconditioning” or “repurposing” irradiated fuel that has reached its burnup limit, in order to fabricate new fuel. These are actually euphemisms for types of reprocessing and would undermine the main purpose of the breed-and-burn reactor: to use uranium more efficiently without any reprocessing.

Proliferation/Terrorism Risk

There are several factors to consider in assessing whether a transition from today’s LWRs to once-through breed-and-burn reactors would have overall benefits for nonproliferation and nuclear security.

First, if once-through breed-and-burn reactors could achieve the benefits of uranium conservation and waste management that are attributed to breeder reactors’ closed fuel cycles, deployment of these reactors could undercut the

rationale for reprocessing and recycling and provide a more secure alternative. This is the strongest selling point of the breed-and-burn concept.

Another advantage is that breed-and-burn reactors would require less enrichment capacity for each unit of electricity generated, on average. TerraPower estimates that the lifetime-averaged separative work for its reactor would be 30,000 separative work units per GWe-yr, about 25 percent of that required for an LWR. If such reactors were to completely replace LWRs, the number and capacity of uranium enrichment plants around the world needed to support the nuclear fuel cycle would decrease, with a potential benefit for nonproliferation. This benefit would be offset somewhat by the fact that the smaller enrichment plants would need to produce ton quantities of HALEU up to just under 20 percent U-235, which, as discussed in chapter 4, is Category II material that raises additional proliferation and security concerns relative to the lower enrichments needed for LWR fuel.

NEED FOR MORE FREQUENT REFUELING

A key issue affecting safeguards is the extent to which, in order to use uranium more efficiently, breed-and-burn reactors would require more frequent refueling or fuel shuffling, possibly even shuffling of small fuel segments in three dimensions. Safeguards for current-generation on-line refueled reactors, such as CANDU reactors, require more inspection resources than LWRs. As discussed in chapters 6 and 7, NLWRs with on-line refueling such as pebble-bed HTGRs would be even more challenging to safeguard than CANDU reactors, given the greater number of items to be tracked and the increased complexity of fuel movements. Item monitoring for safeguards can be difficult if large numbers of items are involved. At the far end of the spectrum, MSRs, which would be continuously refueled and co-located with reprocessing plants, would require safeguards comparable to those at bulk-handling facilities.

Therefore, additional resources and technological advances are needed for on line—refueled reactor safeguards. Ultimately, assurances could be provided with a combination of unannounced and short-notice random inspections, unattended monitoring systems, and inspections to verify the absence of undeclared reprocessing facilities. Such elements are part of the integrated safeguards approach for Canada, a country with CANDU reactors but no reprocessing plants (Whiting, Hosoya, and Doo 2006).

In any event, the TerraPower traveling wave reactor may not be more difficult to safeguard than LWRs since it does not require the reactor vessel to be opened when the

reactor is shut down for fuel reshuffling, and would not normally require safeguards inspections over its lifetime, unless there were unexpected shutdowns. On the other hand, the inability for inspectors to directly verify the amount of fuel in the core could hamper their ability to accurately account for the quantity of fissile material over the reactor's lifetime.

NEED FOR SPENT FUEL "RECONDITIONING"

Another safeguards issue relates to the proposals for technologies for reconditioning breed-and-burn reactor spent fuel to further increase uranium utilization. Although the processes in question might not involve a complete separation of weapon-usable TRU from other materials, they would separate out most of the cesium-137, greatly reducing the self-protecting radiation barrier, and thus could present proliferation and security risks similar to those of conventional reprocessing. In addition, these bulk processes would share some of the same material accountancy difficulties as conventional reprocessing and fuel fabrication plants, reducing the attractiveness of breed-and-burn systems from a non-proliferation perspective. The plutonium content in the discharged fuel is also quite high (more than 12 percent for some designs), which raises safeguards concerns. Although the ability to reuse spent fuel in these systems remains a selling point for some designers, the achievable gains in uranium utilization would not be worth the cost and risk of these processes.

SPENT FUEL DISPOSAL

The ultimate fate of the spent fuel from breed-and-burn reactors is an open question. TerraPower is a metal-fueled fast reactor similar to the EBR-II, and the fuel would be high-burnup by necessity. As discussed earlier, the Department of Energy continues to pyroprocess high-burnup EBR-II spent fuel because it maintains that such fuel cannot be safely disposed of in a repository. The reason for that is the presence of the metallic sodium bond material, which chemically reacts violently with water. However, it appears that sodium bonding cannot be used in breed-and-burn fast reactors because it makes the positive void reactivity effect unacceptably large (Hejzlar et al. 2013). Thus, there would be no justification to pyroprocess the fuel for final disposal. Nevertheless, TerraPower has not yet developed an acceptable substitute for the sodium-bonded fuel. If TerraPower spent fuel has to be reprocessed for disposal, it would nullify the nonproliferation benefits of the breed-and burn concept.

On balance, it appears that the increased risk of diversion of spent fuel at once-through reactors with on-line refueling or shuffling would be a reasonable tradeoff if such reactors could use uranium more efficiently without reprocessing and recycling.

Conclusions and Recommendations

Conclusions

The NLWR landscape is vast and complex, and it is beyond the scope of this report to survey the entire field in depth. Nevertheless, enough is clear even at this stage to draw some general conclusions regarding the safety and security of NLWRs and their prospects for rapid deployment.

Based on the available evidence, it is far from obvious that the NLWR designs currently under consideration, except possibly once-through, breed-and-burn reactors, offer improvements over LWRs significant enough to justify their many risks. Regulators and other policymakers would be wise to look more closely at the nuclear power programs underway to make sure they prioritize safety and security. Future appropriations for NLWR technology research, development, and deployment should be guided by realistic assessments of the likely societal benefits that would result from the investment of billions of taxpayer dollars.

Little evidence supports claims that NLWRs will be significantly safer than today's LWRs. While some NLWR designs offer some safety advantages, all have novel characteristics that could render them less safe.

All NLWR designs introduce new safety issues that will require substantial analysis and testing to fully understand and address—and it may not be possible to resolve them fully. To determine whether any NLWR concept will be significantly safer than LWRs, the reactor must achieve an advanced stage of technical maturity, undergo complete comprehensive safety testing and analysis, and acquire significant operating experience under realistic conditions.

The claim that any nuclear reactor system can “burn” or “consume” nuclear waste is a misleading oversimplification. Reactors can actually use only a fraction of spent nuclear fuel as new fuel, and separating that fraction increases the risks of nuclear proliferation and terrorism.

No nuclear reactor can use spent nuclear fuel directly as fresh fuel. Instead, spent fuel has to be “reprocessed”—chemically treated to extract plutonium and other transuranic elements, which must then be refabricated into new fuel. This introduces a grave danger: plutonium and other transuranic elements can be used in nuclear weapons. Reprocessing and recycling renders these materials vulnerable to diversion or theft and increases the risks of nuclear proliferation and terrorism—risks that are costly to address and that technical and institutional measures cannot fully mitigate. Any fuel cycle that requires reprocessing poses inherently greater proliferation and terrorism risks than the “once-through” cycle with direct disposal of spent fuel in a geologic repository.

Some NLWRs have the potential for greater sustainability than LWRs, but the improvements appear to be too small to justify their proliferation and safety risks.

Although some NLWR systems could use uranium more efficiently and generate smaller quantities of long-lived transuranic isotopes in nuclear waste, for most designs these benefits could only be achieved by repeatedly reprocessing spent fuel to separate out these isotopes and recycle them in new fuel—and that presents unacceptable proliferation and security risks. In addition, reprocessing plants and other associated fuel cycle facilities are costly to build and operate,

and they increase the environmental and safety impacts compared with the LWR once-through cycle. Moreover, the sustainability increases in practice would not be significant in a reasonably foreseeable timeframe.

Once-through, breed-and-burn reactors have the potential to use uranium more efficiently without reprocessing, but many technical challenges remain.

One type of NLWR system that could in principle be more sustainable than the LWR without increasing proliferation and terrorism risks is the once-through, breed-and-burn reactor. Concepts such as TerraPower's traveling-wave reactor could enable the use of depleted uranium waste stockpiles as fuel, which would increase the efficiency of uranium use. Although there is no economic motivation to develop more uranium-efficient reactors at a time when uranium is cheap and abundant, reducing uranium mining may be beneficial for other reasons, and such reactors may be useful for the future. However, many technical challenges would have to be overcome to achieve breed-and-burn operation, including the development of very-high-burnup fuels. The fact that TerraPower suspended its project after more than a decade of development to pursue a more conventional and far less uranium-efficient SFR, the Natrium, suggests that these challenges have proven too great.

High-assay low enriched uranium (HALEU) fuel, which is needed for many NLWR designs, poses higher nuclear proliferation and nuclear terrorism risks than the lower-assay LEU used by the operating LWR fleet.

Many NLWR designs require uranium enriched to higher levels than the 5 percent U-235 typical of light-water reactor fuel. Although uranium enriched to between 10 and 20 percent U-235 (defined here as HALEU) is considered impractical for direct use in nuclear weapons, it is more attractive for weapons use—and requires more stringent security—than the lower-assay enriched uranium in current LWRs.

The significant time and resources needed to safely commercialize any NLWR design should not be underestimated.

It will likely take decades and many billions of dollars to develop and commercially deploy any NLWR design, together with its associated fuel cycle facilities and other support activities. Such development programs would come with a significant risk of delay or failure and require long-term stewardship and funding commitments. And even if a commercially workable design were demonstrated, it would

take many more years after that to deploy a large number of units and operate them safely and reliably.

Vendors that claim their NLWRs could be commercialized much more quickly typically assume that their designs will not require full-scale performance demonstrations and extensive safety testing, which could add well over a decade to the development timeline. However, current designs for sodium-cooled fast reactors and high-temperature gas-cooled reactors differ enough from past reactor demonstrations that they cannot afford to bypass additional full-scale prototype testing before licensing and commercial deployment. Molten salt-fueled reactors have only had small-scale demonstrations and thus are even less mature. NLWRs deployed commercially at premature stages of development run a high risk of poor performance and unexpected safety problems.

Recommendations

The DOE should suspend the Advanced Reactor Demonstration Program pending a finding by the NRC whether it will require prototype testing before licensing the two chosen designs as commercial power reactors.

The DOE has selected two NLWR designs, the Natrium SFR and the Xe-100 pebble-bed HTGR, for demonstration of full-scale commercial operation by 2027. However, the NRC has yet to evaluate whether these designs are mature enough that it can license them without first obtaining data from prototype plants to demonstrate novel safety features, validate computer codes, and qualify new types of fuel in representative environments. Without such an evaluation, the NRC will likely lack the information necessary to ensure safe, secure operation of these reactors. The DOE should suspend the Advanced Reactor Demonstration Program until the NRC—in consultation with the agency's Advisory Committee on Reactor Safeguards and external experts—has determined whether prototypes will be needed first.

Congress should require that an independent, transparent, peer-review panel direct all DOE R&D on new nuclear concepts, including the construction of additional test or demonstration reactors.

Given the long time and high cost required to commercialize NLWR designs, the DOE should provide funding for NLWR R&D judiciously and only for reactor concepts that offer a strong possibility of significantly increasing safety and security—and do not increase proliferation risks. Moreover, unlike the process for selecting the two reactor designs for the

Advanced Reactor Demonstration Program, decisionmaking should be transparent.³¹ Congress should require that the DOE convene an independent, public commission to thoroughly review the technical merits of all NLWR designs proposed for development and demonstration, including those already selected for the ARDP. The commission, whose members should represent a broad range of expertise and perspectives, would recommend funding only for designs that are highly likely to be commercialized successfully while achieving clearly greater safety and security than current-generation LWRs.

The DOE and other agencies should thoroughly assess the implications for proliferation and nuclear terrorism of the greatly expanded production, processing, and transport of high-assay low-enriched uranium (HALEU) required to support the widespread deployment of NLWRs.

Large-scale deployment of NLWRs that use HALEU fuel will require establishing a new industrial infrastructure for producing and transporting the material. The DOE is actively promoting the development of HALEU-fueled reactor designs for export. Given that HALEU is a material of higher security concern than lower-assay LEU, Congress should require that the DOE immediately assess the proliferation and nuclear terrorism implications of transitioning to the widespread use of HALEU worldwide. This assessment should also address the resource requirements for the security and safeguards measures needed to ensure that such a transition can occur without an unacceptable increase in risk.

The United States should make all new reactors and associated fuel facilities eligible for IAEA safeguards and provide that agency with the necessary resources for carrying out verification activities.

The International Atomic Energy Agency, which is responsible for verifying that civilian nuclear facilities around the world are not being misused to produce materials for nuclear weapons, has limited or no experience in safeguarding many types of NLWRs and their associated fuel cycle facilities. NLWR projects being considered for deployment in the United States, such as the Sodium SFR and the Xe-100 pebble-bed HTGR, would provide ideal test beds for the IAEA to develop safeguards approaches. However, as a nuclear-weapon state, the United States is not obligated to give the IAEA access to its nuclear facilities. To set a good example and advance the cause of nonproliferation, the United States should immediately provide the IAEA with permission and funding to apply safeguards on all new US nuclear facilities, beginning at the design phase. This would help to identify safeguard challenges early and give the IAEA experience in verifying similar facilities if they are deployed in other countries.

The DOE AND Congress should consider focusing nuclear energy R&D on improving the safety and security of LWRs, rather than on commercializing immature NLWR designs.

LWR technology benefits from a vast trove of information resulting from many decades of acquiring experimental data, analysis, and operating experience—far more than that available for any NLWR. This gives the LWR a significant advantage over other nuclear technologies. The DOE and Congress should do a more thorough evaluation of the benefits of focusing R&D funding on addressing the outstanding safety, security, and cost issues of LWRs rather than attempting to commercialize less mature reactor concepts. If the objective is to expand nuclear power to help deal with the climate crisis over the next few decades, improving LWRs could be a less risky bet.

Simple Models of Fast Burner/Breeder Cycles

This appendix presents simple models to illustrate the practical ability of fast burner or breeder reactors to meet the sustainability objectives of (1) significantly reducing long-lived, heat-generating transuranic elements (TRU) in radioactive wastes requiring long-term geologic disposal; and (2) significantly increasing natural uranium utilization.

How Long Would It Take to Reduce Transuranics by a Factor of 10 with a Burner Reactor System?

This section estimates the performance of a typical fast burner reactor system in terms of the operating time necessary to achieve a significant reduction in transuranics (TRU) relative to the amount generated by the light-water reactor (LWR) once-through cycle. Note that these estimates have large uncertainties because they are very sensitive to the assumptions going into the calculation, and thus are only illustrative. However, the uncertainties do not affect the general conclusion that fast waste burner systems are not very effective over realistic timescales.

A 1000 megawatt-electric (MWe) LWR produces about 220 kilograms (kg) of TRU in its spent fuel each year, for a total of about 13.2 metric tons after 60 years of operation. How does this compare to the amount of TRU that would remain if a 1000 MWe fast burner reactor had operated instead?

Figures A-1 and A-2 (p. 119) illustrate the fuel cycle for GE-Hitachi's 1000 MWe, PRISM metal-fueled fast reactor, configured as a burner. The TRU burning performance of such a reactor can be characterized by the conversion ratio. This reactor has a conversion ratio of 0.5, which means that in each operating cycle, one half as many TRU isotopes are produced as are fissioned. This is an unrealistically low conversion ratio that was chosen to maximize the TRU-burning ability of the reactor in this calculation. In practice, a fast reactor with such a low conversion ratio would be particularly challenging to operate and less safe than fast reactors with higher conversion ratios (Hoffman, Yang, and Hill 2006.)

Figure A-1 shows how the TRU fuel needed for the start-up of the fast burner reactor is obtained from reprocessing

the spent fuel from LWRs. An amount of spent fuel equivalent to that generated from one year's operation of 67 1,000-MWe LWRs (less than 2 percent of the current US spent fuel stockpile) would need to be reprocessed to obtain the TRU fuel for the first core and the first full reload core for the fast burner reactor. Also note that more than 12,000 metric tons of natural uranium was enriched to produce the low-enriched uranium fuel to supply the 67 LWRs. Enriching this uranium also generated a stockpile of more than 11,000 metric tons of depleted uranium waste.


Figure A-2 shows the fueling requirement for the fast burner reactor after it reaches an equilibrium state (meaning that the fuel requirement remains constant over time). At equilibrium, 42.7 metric tons of LWR spent fuel would be needed per year, corresponding to the discharge from two LWRs (or about 2 percent of current annual US spent fuel generation), to provide makeup fuel to replace the TRU fissioned in the fast reactor.

Therefore, the US spent fuel stockpile and current rate of spent fuel generation could fuel 50 1,000-MWe PRISM fast burner reactors. At first glance, this appears to say that if one PRISM reactor were built for every two operating LWRs in the United States, the PRISM reactors would be able to consume all US spent fuel, as promoters of the technology claim. But the more important consideration in terms of waste burning capability is not the mass of TRU that are fed into the system, but what would remain when the first generation of facilities is shut down—at which time society would have to decide whether to replace them or to dispose of the left-over material.

After 60 years of operation, the quantity of TRU in one reactor core and its corresponding fuel cycle facilities would be about 11.7 metric tons, assuming that each facility stores one year's worth of material. Compare this quantity to the TRU that would need to be disposed of if an LWR had operated over this time instead of the fast burner reactor. As discussed above, a 1000 MWe LWR would produce about 13.2 metric tons after 60 years of operation. But the 41 metric tons of TRU that would have been used as fuel for the burner reactor

continued on p. 120


FIGURE A-1. Heavy Metal Mass Flow for Startup and Transition Core for a Fast Burner Reactor



Note: Assumes LWRs will only need to supply one full core reload before fast reactor becomes self-sustaining. Zero process loss assumed.

SOURCE: ADAPTED FROM HOFFMAN, YANG, AND HILL 2006.

FIGURE A-2. Annual Heavy Metal Mass Flow at Equilibrium for a Fast Burner Reactor



Note: Zero processed loss assumed.

SOURCE: ADAPTED FROM HOFFMAN, YANG, AND HILL 2006.

continued from p. 118

also has to be counted. Thus, the total quantity of TRU that would have to be disposed of would be 13.2 metric tons plus 41 metric tons, for a total of 54.2 metric tons.³² The resulting TRU reduction factor would be 4.4, well below the Department of Energy's (DOE) factor-of-10 standard (see chapter 3). If the reactor were to shut down at that point, the unburned TRU inventory would have to be either disposed of in a geologic repository or used in a replacement burner reactor. Even after another 60-year run, the reduction would be only about 8.8—still below the DOE criterion.

The above model is unrealistic in that it assumes that the fast burner reactor's spent fuel can be immediately reprocessed and recycled into new fuel. In practice, the spent fuel would have to be stored for several years at a minimum until it had cooled sufficiently to be safely reprocessed. This would increase the quantity of TRU in the fuel cycle that would have to be counted after the system shuts down. If one assumes a three-year cooling period in the above example, then at least one core's worth of spent fuel would be stored at any time during the equilibrium cycle. Thus, at shutdown after 60 years, the unfissioned TRU in the reactor and fuel cycle would total around 20 metric tons—and the TRU reduction factor would only be 2.7. If the required cooling period were longer, or if extra TRU had to be stored at the fuel fabrication facility, the reduction factor would be even less favorable. This example illustrates how sensitive the performance of a TRU burner system is to changes in the assumptions.

IMPORTANCE OF PROCESS LOSSES

The above model is unrealistic also because it does not account for process losses—TRU that ends up in the waste streams of the reprocessing and fuel fabrication plants and is not recycled. Process losses will never be zero in any realistic system because it is not feasible or economical to recover every TRU atom from waste streams.

The importance of process losses to the performance of a TRU burner system can be seen as follows. Considering the fast-burner reactor cycle in Figure A-2, assume that the quantity of TRU lost to waste streams amounts to 2 percent of each facility's throughputs, which is the lower limit of the observed process losses from metal fast reactor fuel pyroprocessing and fuel fabrication operations (Westphal et al. 2017; Hayes 2017). (The upper limit is much higher—over 25 percent for metal fuel fabrication alone (Moore and Severynse 2020).) Then, about 87 kg of TRU (and 180 kg of uranium) would end up as waste every year (Figure A-2, p. 119). In order to compensate for this loss, an additional eight metric tons of LWR

spent fuel would have to be reprocessed to supply the additional TRU for fuel.

At the end of the 60-year reactor lifetime, 5.2 metric tons of TRU would be contained in the processing waste. Adding this to the TRU inventory in the reactor and fuel cycle for the no-cooling scenario, the TRU reduction factor would be $(54.2 + 5.2)/(12.4 + 5.2) = 3.4$, or 25 percent smaller than the reduction factor of a system with zero process losses. In this case, it would take 177 years to reach a factor-of-10 TRU reduction.

Thus, the system of reprocessing plants, fuel fabrication plants, fast reactors, and associated facilities would have to operate over a period spanning many generations—and be rebuilt many times—before it could achieve a significant reduction in TRU mass and a significant benefit for geologic disposal. The minor benefits of running such a system for a shorter period of time—one or two facility lifetimes—would not justify the cost premium required for development, deployment, and operation of a fleet of fast reactors and associated fuel cycle facilities compared to LWRs on a once-through cycle (see chapter 5).

How Long Would It Take for a Fast Breeder Reactor to Extract 100 Times As Much Energy from Uranium Ore as an LWR?


Figures A-3 and A-4 (p. 121) illustrate the fuel cycle for GE-Hitachi's 1000 MWe, PRISM-type metal-fueled fast reactor configured as a breeder, rather than a burner as in the above example. Rather than having a conversion ratio of less than 1 as in the previous example, it has a breeding ratio of greater than 1, indicating that more TRU isotopes are produced during each cycle than are fissioned.

As shown in Figure A-3, 17.3 metric tons of plutonium and other TRU elements would be needed as driver fuel for the first core and the first reloads of a 1000 MWe PRISM-type fast breeder reactor (Dubberly, Wu, and Kubo 2003).³³ This would require reprocessing 1600 metric tons of spent fuel from LWRs, corresponding to the annual discharge of 80 1,000-MWe LWRs. About 14,750 metric tons of natural uranium was mined to supply the LWRs, and enrichment of the natural uranium generated a stockpile of more than 13,000 metric tons of depleted uranium. (The TRU in the current US stockpile of spent fuel could be used to start up about 50 such reactors).

Once started, the reactor would need to operate for a number of cycles to reach an equilibrium state. At that point, it would only need a fresh supply of about 1.1 metric tons of

continued on p. 122


FIGURE A-3. Production of Initial Core and First Reload for a Fast Breeder Reactor



Notes: Assumes LWRs will only need to supply one full core reload before fast reactor becomes self-sustaining. Zero process loss assumed.

SOURCE: ADAPTED FROM DUBBERLY, WU, AND KUBO, 2003.

FIGURE A-4. Annual Heavy Metal Mass Flow at Equilibrium for a Fast Breeder Reactor



Note: Zero processed loss assumed.

SOURCE: ADAPTED FROM DUBBERLY, WU, AND KUBO, 2003.

continued from p. 120

depleted uranium blanket fuel from the stockpile each year (Figure A-4, p. 121). But the reactor cannot operate on blanket fuel alone. Plutonium and other TRU separated from the fast reactor spent fuel by reprocessing would be needed to replenish the driver fuel.

Over a 60-year lifetime at an 85 percent capacity factor, the fast breeder reactor would fission about 50 metric tons of heavy metal. Thus the uranium utilization of the breeder system would be about 0.3 percent of the 14,750 metric tons of natural uranium ore that was used to fuel the LWRs that generated the TRU for the startup core and first reloads. Since, however, the same quantity of uranium ore would have previously been used to fuel 80 LWRs for one year, the uranium utilization efficiency of the fast breeder plus the LWR system after 60 years of fast reactor operation would be around 0.6 percent + 0.3 percent = 0.9 percent. This is only a minor (50 percent) increase in uranium utilization over that of LWRs alone. In order for the fast breeder to achieve 100 times the uranium utilization of LWRs alone—that is, 60 percent—the system would have to operate for around 11,000 years. And if process losses are taken into account, more uranium would have to be enriched and irradiated to produce the same quantity of TRU in fabricated fuel, and uranium utilization would be even lower.

Uranium Utilization Efficiency of Burner Reactors vs. Breeder Reactors

Figure A-2 illustrates why a fast reactor designed to burn the TRU in LWR spent fuel would not utilize uranium very efficiently, ultimately achieving only a 33 percent increase over the LWRs alone. As shown in the figure, every year the burner reactor is fed the TRU from 42.7 metric tons of spent fuel—corresponding to the annual fuel discharge from 2.1 1000-MWe LWRs (at 20 metric tons per GWe-year). The quantity of uranium ore used to produce that fuel is about 380 metric tons. The amount of heavy metal the combined LWR-fast burner system would fission every year would be about 3 metric tons; therefore, the uranium utilization efficiency would be $3/380 = 0.8$ percent, compared to 0.6 percent for the LWRs.

Waste Reduction Factor of a Breeder Reactor

Similarly, if one operates a fast reactor system as a breeder reactor, it will not be effective in reducing the quantity of long-lived TRU that would require disposal in a geologic repository.

How does the quantity of TRU discarded to waste by a breeder reactor system compare to that generated by an LWR?

First, consider only the steady-state production of waste TRU that would be generated each year. As discussed earlier, a 1000 MWe LWR discharges about 220 kg of TRU in its spent fuel each year. For the fast breeder reactor cycle, most of the TRU in the spent fuel is recycled, but each time that exceeds 10 percent, which is more realistic,

Consider the process losses for the fast breeder reactor cycle in Figure A-4. Assume, as in the earlier example, that the quantity of TRU lost to waste streams amount to 2 percent of each facility's throughputs. Then, from Figure A-2, about 60 kg of plutonium and other TRU (and more than 500 kg of uranium) would end up as waste every year. (The TRU process loss could be made up for by a portion of the excess 230 kg of TRU bred every year, reducing the effective breeding ratio and the excess TRU amount to 170 kg.) Thus, even if the material in the system is not taken into account, the quantity of TRU that would need to be sent to a repository would be reduced only by a factor of four (60 kg/220 kg) compared to the LWR, a reduction that does not meet the DOE's factor-of-10 criterion. And if the process loss approaches 10 percent, the TRU in the breeder waste stream could actually exceed the amount discharged by LWRs in a once-through cycle.

To compare apples to apples, a more rigorous calculation would also account for the quantities of the TRU in the cores and fuel cycles of both the LWR and the LWR plus fast breeder systems. From Figure A-3, it can be seen that one 1000 MWe LWR would need to run for 80 years to provide the startup TRU for one 1000 MWe fast breeder reactor. What is the net TRU waste generated by this LWR and fast breeder system?

By the end of the 60-year fast breeder reactor lifetime, the system would have generated 140,000 MWe-years of electricity (80,000 in the LWR and 60,000 in the fast breeder). The remaining TRU to be disposed of would include, in addition to the 3.6 metric tons of TRU in the processing waste, 8.7 metric tons of TRU in the reactor core and 2.8 metric tons of TRU in fuel fabrication and reprocessing plants servicing the reactor. (This does not include the 170 kg of excess TRU that was bred over the reactor's lifetime, which presumably has been used to start up another fast reactor and thus does not count as waste.) This unburned TRU would either have to be disposed of in a geologic repository or used in a replacement breeder reactor. If the next generation of humans decides not to replace this fast breeder reactor in 60 years, the total amount of plutonium to be disposed of would include the material in the system—which is 11.5 metric tons. Thus, including the TRU waste from process losses, more than 15 metric tons of TRU in the fast breeder core and fuel

cycle would require disposal. The core of the shutdown LWR would also contain about 0.5 metric ton of TRU requiring disposal, for a total of about 15.5 metric tons.

This figure should be compared to the TRU that would have to be disposed of after 140,000 MWe of LWR generation, which would be 140 x 220 kg plus the 0.5 metric ton in the reactor core, for a total of about 31 metric tons of TRU. Thus, the TRU reduction factor would be $31/15.5 = 2$, also far below the DOE's criterion.

If the next generation did decide to build a replacement breeder reactor and operate the system for another 60 years, there would be 7.2 metric tons of TRU in the waste stream but still 11.5 metric tons in the reactor and fuel cycle, for a total of 18.7 metric tons of TRU. The TRU reduction factor would only be slightly better at 2.35. Again, this is far below the DOE's factor-of-10 criterion. In this model, analysis shows that a factor-of-10 TRU reduction would be impossible to achieve, no matter how long the system operated.

The assumed size of the process losses is a critical factor in this analysis, since they represent TRU that is lost from the system and cannot be burned. The DOE often assumes a much lower process loss fraction of 0.1 percent, which is hundreds of times smaller than what has been achieved for pyroprocessing and metal fuel fabrication in practice. But even with such a low process loss, the leftover TRU from the breeder reactor after 60 years would be 11.7 metric tons, resulting in a reduction factor of 2.65 instead of 2. This idealized system would be able to meet the DOE's factor of 10 reduction eventually, but it would have to operate for over 500 years—nine reactor lifetimes.

These examples illustrate that in practice a reactor and fuel cycle system configured to increase the efficient use of uranium by breeding TRU cannot effectively burn TRU, and vice-versa.

[ENDNOTES]

- 1 About 10 percent of reactors today use heavy water (in which hydrogen is replaced by the isotope deuterium) as a coolant.
- 2 The projected nuclear power capacity in 2050 in the IAEA's 2020 high case is about four percent lower than its 2018 projection and 18 percent lower than its 2017 projection, illustrating the variability in such forecasts.
- 3 The think tank Third Way has identified 71 "advanced" nuclear power development projects in North America, of which nearly all are NLWRs (Milko, Allen and Fitzpatrick 2018). Of those, 64 are private sector projects, with the others housed at US national laboratories and universities that receive government funding. This list also includes many nuclear fusion reactor projects, which are further away from commercialization than fission-based reactors. The actual number of fission reactor ventures with significant private funding is relatively small. See, for example, Morgan et al. 2018.
- 4 This report addresses only NLWRs, and not small modular LWRs such as NuScale or other novel LWR concepts such as the supercritical light-water reactor. The Union of Concerned Scientists previously evaluated issues related to small modular LWRs in its report *Small Isn't Always Beautiful* (Lyman 2013). The present report also does not discuss nuclear fusion reactors.
- 5 Some analysts have rightfully questioned the effectiveness of the DOE's reactor research and development programs for new nuclear reactor technologies (Ford et al. 2017).
- 6 NLWRs are not the only reactor options for providing high-temperature process heat. A type of light-water reactor called the "supercritical light-water reactor" could also operate at high temperature. While operating experience from current-generation LWRs could give the supercritical LWR an advantage over NLWRs, the design still would require significant research and development. Although the DOE stopped funding supercritical LWR research years ago, other countries continue to pursue it. Also, it is possible to amplify the outlet temperature of a LWR. In addition, the DOE is also pursuing the production of hydrogen fuel—one important use of nuclear process heat—at currently operating LWRs. Thus higher-temperature reactors are not essential for this application, although they may be preferred.
- 7 This estimate assumes that the feed material for the demonstration plant will be 4.95 percent-enriched uranium purchased on the open market (Dyke 2020). However, it is unclear whether a supply of enriched feed could be found that would be entirely US-origin. The DOE says that production of US-origin fuel is one of the key reasons for its support of this project. If natural uranium must be used, then the facility could only produce about 130 kilograms of 19.75 percent-enriched HALEU per year, based on company information (Dyke 2020).
- 8 There is currently no consistent definition of HALEU in the literature. Some sources define it as the entire range of enrichments between 5 and 20 percent. This report adopts the working definition introduced by URENCO-USA for LEU+ and HALEU, which directly corresponds to their different security requirements (Fletcher 2020).
- 9 Oxide spent fuels can only be pyroprocessed if they are first "reduced"—chemically treated to remove oxygen and convert them to a metal form.
- 10 The Global Nuclear Energy Partnership's Nonproliferation Impact Assessment was never issued in final form because President Obama cancelled the program soon after he took office in January 2009.
- 11 For micro-reactors such as the Aurora, the engineering demonstration and performance demonstration steps would be the same since scale-up would not be necessary.
- 12 In 1954, prior to the EBR-1 accident, Argonne National Laboratory conducted a series of tests in Idaho called BORAX in which reactor cores were deliberately destroyed.
- 13 The Nuclear Energy Agency's fuel cycle calculations assume that the fleet consists of 82.5 percent LWRs and 17.5 percent fast reactors. It uses a fast reactor cost premium of 20 percent but also assumes that the cost of non-reactor facilities are the same for processing plutonium fuel for both thermal and fast reactors, based on the limited share of fast reactors in the mix. This assumption is unrealistic with respect to both safety and security costs.
- 14 There are differing opinions on how to interpret this statement, and the NRC is currently considering a proposal to relax the criterion for NLWRs by changing the language to "the reactor core and associated systems that contribute to reactivity feedback" (emphasis added). This change would give credit to factors other than prompt feedback that could reduce reactivity. It would remove a potential obstacle to the licensing of fast reactors or other systems that have troublesome inherent instabilities and may need to rely on other feedback effects, such as expansion of core structures, to provide negative reactivity.
- 15 This burnup may actually be too high. The DOE has limited the average discharge burnup for PRISM-type metal fuel for the very similar Versatile Test Reactor to about 6 percent (approximately 60,000 MWd/MTHM) in order to remain within the envelope of the current fuel qualification database (Youinou et al. 2020).
- 16 This is different from the Very High Temperature Reactor, or VHTR, a less mature design originally included in the Generation IV program that would have an outlet temperature up to 950°C. New materials would need to be developed and qualified for performance at such high temperatures (Petti et al. 2017).
- 17 The 2017 DOE advanced demonstration and test reactor study considers Peach Bottom to be an engineering demonstration and the latter two to be performance demonstrations (and not commercial demonstrations) for current-design HTGRs, which would be less than 300 MWe and have different features such as passive safety systems (Petti et al. 2017).
- 18 While China refers to the HTR-10 as a test reactor, the 2017 DOE demonstration and test reactor study categorizes it as an engineering-scale demonstration reactor (Petti et al. 2017).
- 19 One may reasonably wonder why a reactor called the Xe-100 would have a generating capacity of 200 MWth. The answer is that the original reactor concept was only 100 MWth (DOE-NE 2014). As is the case with other small modular reactors, including the NuScale small modular reactor (SMR), designers have steadily increased the power ratings of each module, presumably to improve the economics. Also, the reported electrical generating capacity of each module recently increased from 75 MWe to 80 MWe.
- 20 General Atomics is also developing "accident tolerant" fuels for operating LWRs that are based on silicon carbide technology, in an example of technology transfer from HTGRs to LWRs.
- 21 Because of the generally poor economics of small modular reactors (Lyman 2013), this limitation could pose a problem for the eventual commercial viability of HTGRs.
- 22 This reactor may actually not be a pebble-bed design, given that a reactor requiring continuous refueling likely would not be suitable for this application.
- 23 A molten core draining mechanism is not a unique feature of MSR but can also be incorporated into LWRs. The EPR under construction at sites in Finland and France, is a Generation III LWR with a "core catcher." In the event of a Fukushima-like accident where the fuel melts through the bottom of the reactor vessel, the EPR is designed so that the molten core will flow into a chamber below the containment floor where it can be safely cooled. Unlike the MSR, though, this is a last-resort measure in the unlikely event of a core melt. For the MSR, core draining is the first line of defense.
- 24 The data in Gehin and Powers (2016) were adjusted here by using a uranium tails assay of 0.25 percent, to be consistent with other calculations in the present report.
- 25 The MSBR TRU waste quantity here is about twice the value quoted in Table V of Gehin and Powers (2016) because it also includes the end-of-life MSBR core inventory, which the study did not consider as waste material.
- 26 These numbers reflect that in recent documents Terrestrial Energy has updated the IMSR400 from its original capacity of 400 MWth (190 MWe). In addition, the anticipated thermal efficiency has apparently decreased from over 48 percent to 43 percent or less.
- 27 To arrive at this figure, one must adjust the results in Figure 9 of Betzler, Powers, and Worrall (2017), which presents the results of a two-dimensional simulation, by adding the additional plutonium predicted by the three-dimensional simulation in Figure 43.

- 28 The Oak Ridge review found that the Transatomic Power reactor's spent fuel would contain only about half of the total amount of "actinide waste" that a comparable LWR would generate (Robertson et al. 2017). However, the category of "actinide waste" includes uranium as well as the long-lived transuranic elements, and excluding the uranium in the spent fuel paints a different picture.
- 29 The reason is that the Transatomic Power reactor is a spectral shift reactor that uses movable moderator rods to control the neutron spectrum over the reactor's lifetime. In the early years of operation, the neutron spectrum is relatively hard, which results in a rapid buildup of plutonium. The plutonium inventory in the core rises to nearly four metric tons. Subsequently, the spectrum is softened, which promotes fission of the accumulated plutonium, but the in-core inventory remains high.
- 30 As of early 2021, Terrestrial Energy has said that final design of its off-gas system is still under development, but "presently the off-gas system is NOT used as a system that would clean up containment airborne radionuclides" (Terrestrial Energy 2021).
- 31 Although the DOE has said that an external review of its selections took place, it has not publicly released the reviewers' names and affiliations—nor has it publicly documented their findings.
- 32 The quantity of TRU fed into the burner reactor cycle over 60 years is the quantity in the initial core and first reloads (which last for four years) plus 56 times the annual TRU feed requirement: 14.6 metric tons + 56 x 0.47 metric ton = 41 metric tons of TRU.
- 33 The additional reloads are necessary because the fast reactor will require fissile makeup until it reaches equilibrium (and achieves either break-even or breeding). Studies by GE assume that it will take about a decade of operation for its PRISM reactor to reach equilibrium (Dubberly, Wu, and Kubo 2003); a shorter period is assumed here for ease of illustration.

[REFERENCES]

- ACRS (Advisory Committee on Reactor Safeguards). 2011. *Official Transcript of Proceedings: Advisory Committee on Reactor Safeguards Future Plant Design Subcommittee*. Rockville, Maryland: NRC (Nuclear Regulatory Commission). <https://www.nrc.gov/docs/ML1111/ML111110657.pdf>.
- Adamov, E.O., R.M. Aleksakhin, A.V. Bychkov, A.V. Dedul, Y.S. Mochalov, V.A. Pershukov, A.V. Petrenko, V.I. Rachkov, and D.A. Tolstouhov. 2016. “«PRORYV» Project as Implementation of INPRO Basic Principles for Sustainable Development on the Basis of New Platform for Nuclear Power.” Presented at the IAEA 12th INPRO Dialogue Forum, Vienna, Austria, April 13. https://www-legacy.iaea.org/INPRO/12th_Dialogue_Forum/017_A...Bychkov_Proryv_project.pdf.
- AEC (Atomic Energy Commission). 1972. “An Evaluation of the Molten Salt Breeder Reactor.” Washington, DC. <https://energy-fromthorium.com/pdf/WASH-1222.pdf>.
- Allibert, Michel, Elsa Merle, Sylvie Delpech, Delphine Gerardin, Daniel Heuer, Axel Laureau, and Simon Moreau. 2020. “Preliminary Proliferation Study of the Molten Salt Fast Reactor.” *EPJ Nuclear Sciences & Technologies* 6: 5. <https://doi.org/https://doi.org/10.1051/epjn/2019062>.
- Amano, Yukiya. 2008. “Director General’s Remarks at International Symposium on Uranium Raw Material for the Nuclear Fuel Cycle: URAM-2018.” Speech presented at the International Symposium on Uranium Raw Material for the Nuclear Fuel Cycle, Vienna, Austria, June 25. <https://www.iaea.org/newscenter/statements/director-generals-remarks-at-international-symposium-on-uranium-raw-material-for-the-nuclear-fuel-cycle-uram-2018>.
- ANL (Argonne National Laboratory). 2018. “Recycling Used Nuclear Fuel for a Sustainable Energy Future. Pyroprocessing Technologies.” Department of Energy. https://www.ne.anl.gov/pdfs/Pyroprocessing_brochure_2018.pdf.
- ANS (American Nuclear Society). 2005. “Fast Reactor Technology: A Path to Long-Term Energy Sustainability.” Position Statement 74. La Grange Park, IL. <http://cdn.ans.org/pi/ps/docs/ps74.pdf>.
- AREVA NP Inc. 2010. “Pebble Bed Reactor Technology Readiness Study: Technical Data Record.” 12-9151714-000. <https://art.inl.gov/NGNP/NEAC%202010/Pebble%20Bed%20Reactor%20Technology%20Readiness%20Study%20-%20AREVA.pdf>.
- Back, Christina. 2017. *Hearing on S.512 Nuclear Energy Innovation and Modernization Act*. https://www.epw.senate.gov/public/_cache/files/e/0/e0095898-71fc-4185-b0b0-96e202e4c97b/EE7AA0E9E0E906822FBF1872FEA93771.back-testimony.pdf.
- Bailly, Frederic. 2014. “Progress in Dealing with the World’s Stockpile of Used Nuclear Fuel – How Can Consolidated Interim Storage and Reprocessing Help?” Presented at the Waste Management 2014, Arizona, March. <http://archive.wmsym.org/2014/presentations/361.pdf>.
- Baranwal, Dr. Rita. 2019. “Analysis: FY2020 Spending Bill Points to Nuclear Resurgence.” Energy.Gov. December 2019. <https://www.energy.gov/ne/articles/analysis-fy2020-spending-bill-points-nuclear-resurgence>.
- Bays, Samuel, and Steve Piet. 2010. “The Effect of Burnup and Separation Efficiency on Uranium Utilization and Radiotoxicity.” INL/CON-10-20154. Idaho Falls, ID: Idaho National Laboratory. <https://inldigitallibrary.inl.gov/sites/sti/sti/4731774.pdf>.
- Beck, J M, and L F Pincock. 2011. “High Temperature Gas-Cooled Reactors Lessons Learned Applicable to the Next Generation Nuclear Plant.” INL/EXT-10-19329. Idaho Falls, ID: Idaho National Laboratory. <https://doi.org/10.2172/1023461>.
- Betzler, Powers, Jeffrey J., and Worrall, Andrew. 2017. “Molten Salt Reactor Neutronics and Fuel Cycle Modeling and Simulation with SCALE.” *Annals of Nuclear Energy* 101: 489–503. <https://doi.org/10.1016/j.anucene.2016.11.040>.
- Bowers, Harlan. 2017. “X-Energy and the Xe-100.” Presented at the DC ANS April Meeting, April 5. http://local.ans.org/dc/wp-content/uploads/2014/01/ANS_Xe-100-Overview_04052017.pdf.
- . 2018. “X-Energy: Path to Deployment.” Presented at the Enabling Advanced Reactors for the Market Symposium, George Washington University, March. <https://gain.inl.gov/SiteAssets/EARMSPresentations/08Bowers-X-EnergyARSymposium.pdf>.
- Braun, Chaim. 2012. “Economic Analysis of Fast Spectrum Reactors.” In *Fast Spectrum Reactors*, edited by Alan E. Waltar, Donald R. Todd, and Pavel V. Tsvetkov, 39–46. Boston, MA: Springer US. https://doi.org/10.1007/978-1-4419-9572-8_3.
- Brovchenko, M., D. Heuer, E. Merle-Lucotte, M. Allibert, V. Ghetta, A. Laureau, and P. Rubiolo. 2013. “Design-Related Studies for the Preliminary Safety Assessment of the Molten Salt Fast Reactor.” *Nuclear Science and Engineering* 175 (3): 329–39. <https://doi.org/10.13182/NSE12-70>.
- Brovchenko, Mariya, Jan-Leen Kloosterman, Lelio Luzzi, Elsa Merle, Daniel Heuer, Axel Laureau, Olga Feynberg, et al. 2019. “Neutronic Benchmark of the Molten Salt Fast Reactor in the Frame of the EVOL and MARS Collaborative Projects.” *EPJ Nuclear Sciences & Technologies* 5: 2. <https://doi.org/10.1051/epjn/2018052>.
- Bunn, Matthew, John P. Holdren, Steve Fetter, and Bob Van Der Zwaan. 2005. “The Economics of Reprocessing versus Direct Disposal of Spent Nuclear Fuel.” *Nuclear Technology* 150 (3): 209–30. <https://doi.org/10.13182/NT05-A3618>.
- Busby, Jeremy, Lauren Garrison, Lianshan Lin, Stephen Raiman, Sam Sham, Chinthaka Silva, and Hong Wang. 2019. “Technical Gap Assessment for Materials and Component Integrity Issues for Molten Salt Reactors.” Oak Ridge, TN: Oak Ridge National Laboratory. <https://www.nrc.gov/docs/ML1907/ML19077A137.pdf>.
- Carlson, Donald E., and Sydney J. Ball. 2016. “Perspectives on Understanding and Verifying the Safety Terrain of Modular High Temperature Gas-Cooled Reactors. Proceedings of the HTR 2014, Weihai, China.” *Nuclear Engineering and Design* 306: 117–23. <https://doi.org/10.1016/j.nucengdes.2016.01.015>.
- Chang, Yoon. 1992. “Progress and Status of the Integral Fast Reactor (IFR) Development Program.” In *Proceedings of the American Power Conference*. Chicago, IL: DOE. https://inis.iaea.org/collection/NCLCollectionStore/_Public/23/049/23049227.pdf?r=1&r=1.

- Chiba, Satoshi, Toshio Wakabayashi, Yoshiaki Tachi, Naoyuki Takaki, Atsunori Terashima, Shin Okumura, and Tadashi Yoshida. 2017. "Method to Reduce Long-Lived Fission Products by Nuclear Transmutations with Fast Spectrum Reactors." *Scientific Reports* 7 (1): 13961. <https://doi.org/10.1038/s41598-017-14319-7>.
- Choe, J, M Ivanova, D LeBlanc, S Mohaptra, and R Robinson. 2018. "Fuel Cycle Flexibility of Terrestrial Energy's Integral Molten Salt Reactor (IMSR®)." In *Proceedings of the 38th Annual Conference of the Canadian Nuclear Society and 42nd Annual CNS/CNA Student Conference*. Saskatoon, SK, Canada: Terrestrial Energy Inc. <https://www.terrestrialenergy.com/wp-content/uploads/2018/09/TerrestrialEnergyPaperCNS2018PDF.pdf>.
- Chung, Kiwhan, David H. Beddingfield, William H. Geist, and Sang-Yoon Lee. 2012. "Development of a Reliable Fuel Depletion Methodology for the HTR-10 Spent Fuel Analysis." LA-UR-12-22447. Los Alamos, NM: Los Alamos National Lab. <https://www.osti.gov/biblio/1045380>.
- Cipiti, Benjamin B., Felicia Angelica Duran, Bradley Robert Key, Yaxi Liu, Ivan Lozano, and Rebecca Ward. 2012. "Modeling and Design of Integrated Safeguards and Security for an Electrochemical Reprocessing Facility." SAND2012-9303, 1055891. Albuquerque, NM: Sandia National Laboratories. <https://doi.org/10.2172/1055891>.
- Cipiti, Benjamin B., and Nathan Shoman. 2018. "Bulk Handling Facility Modeling and Simulation for Safeguards Analysis." *Science and Technology of Nuclear Installations* 2018 (October). <https://doi.org/10.1155/2018/3967621>.
- Cipiti, Benjamin B., Nathan Shoman, and Phillip Honnold. Forthcoming. "Safeguards Modeling for Advanced Nuclear Facility Design." *Journal of Nuclear Materials Management*.
- Clemmer, Steve, Jeremy Richardson, Sandra Sattler, and Dave Lochbaum. 2018. "The Nuclear Power Dilemma: Declining Profits, Plant Closures, and the Threat of Rising Carbon Emissions." Cambridge, MA: Union of Concerned Scientists. <https://www.ucsusa.org/resources/nuclear-power-dilemma>.
- Clinton, W. *Nonproliferation and Export Control Policy*. 1993. Presidential Decision Directive.
- Coble, Jamie B., Steven E. Skutnik, S. Nathan Gilliam, and Michael P. Cooper. 2020. "Review of Candidate Techniques for Material Accountancy Measurements in Electrochemical Separations Facilities." *Nuclear Technology* 206 (12): 1803–26. <https://doi.org/10.1080/00295450.2020.1724728>.
- Cochran, Thomas B, Harold A. Feiveson, Walt Patterson, Gennadi Pshakin, M.V. Ramana, Mycle Schneider, Tatsujiro Suzuki, and Frank von Hippel. 2010. *Fast Breeder Reactor Programs: History and Status*. Research Report 8. Princeton, NJ: International Panel on Fissile Materials. <http://libweb.iaea.org/library/eBooks/fast-breeder-reactor.pdf>.
- Crawford, Douglas C. 2020. "VTR Fuel Cycle and Waste Management." Presented at the National Academy of Sciences Committee on the Merits and Viability of Different Nuclear Fuel Cycles and Technology Options and the Waste Aspects of Advanced Nuclear Reactors, December. <https://www.nationalacademies.org/event/12-07-2020/merits-and-viability-of-different-nuclear-fuel-cycles-and-technology-options-and-the-waste-aspects-of-advanced-nuclear-december-7-8-2020-meeting>.
- Dalton, David. 2013. "China Begins Construction Of First Generation IV HTR-PM Unit." The Independent Global Nuclear News Agency. January 2013. <https://www.nucnet.org/news/china-begins-construction-of-first-generation-iv-htr-pm-unit>.
- Demkowicz, Paul, David Petti, and Hans Gougar. 2017. "HTR (TRISO) Fuel Safety Research Activities and Needs." Presented at the GAIN Fuel Safety Research Workshop, Idaho National Laboratory, May 1. <https://gain.inl.gov/SiteAssets/Fuel%20Safety%20Presentations/16%20Demkowicz%20HTGR%20Gap%20Analysis.pdf>.
- DeWitte, Jacob. 2016. *Hearing on the Status of Advanced Nuclear Technologies*. <https://www.energy.senate.gov/services/files/3BA637BA-6A6C-4DF9-9C7F-A0A2BBF8D8E0>.
- Diego Di Sanzo, Christian. 2014. "Increasing Fuel Utilization of Breed and Burn Reactors." PhD, University of California, Berkeley. https://digitalassets.lib.berkeley.edu/etd/ucb/text/DiSanzo_berkeley_0028E_14354.pdf.
- DOD (Department of Defense). 2019. "Pele Request for Solutions Q&A." <https://govtribe.com/file/government-file/requestforsolutionspele-pele-rfs-qas-dot-xlsx>.
- . 2020. "DOD Awards Contracts for Development of a Mobile Microreactor." Press Release. U.S. DEPARTMENT OF DEFENSE. March 9, 2020. <https://www.defense.gov/Newsroom/Releases/Release/Article/2105863/dod-awards-contracts-for-development-of-a-mobile-microreactor>.
- DOE (Department of Energy). 2006. "GNEP Element: Demonstrate More Proliferation-Resistant Recycling." https://www.energy.gov/sites/prod/files/edg/media/GNEP/06-GA50035e_2-col.pdf.
- . 2008. "Draft Global Nuclear Energy Partnership Programmatic Environmental Impact Statement." DOE/EIS-0396. <https://www.energy.gov/nepa/downloads/eis-0396-draft-programmatic-environmental-impact-statement>.
- . 2016a. "Energy Department Announces New Investments in Advanced Nuclear Power Reactors to Help Meet America's Carbon Emission Reduction Goal." EnergyGov. January 15, 2016. <https://www.energy.gov/articles/energy-department-announces-new-investments-advanced-nuclear-power-reactors-help-meet>.
- . 2016b. "Energy Department Announces New Investments in Advanced Nuclear Power Reactors to Help Meet America's Carbon Emission Reduction Goals." Gateway for Accelerated Innovation in Nuclear. January 16, 2016. <https://www.energy.gov/articles/energy-department-announces-new-investments-advanced-nuclear-power-reactors-help-meet>.
- DOE-NE (Department of Energy Office of Nuclear Energy). 2014. "Advanced Reactor Concepts Technical Review Panel, Evaluation and Recommendations for Future R&D on Seven Advanced Reactor Concepts, Conducted March through June 2014." Department of Energy (DOE). <https://www.energy.gov/sites/prod/files/2014/12/f19/Advance%20Reactor%20Concepts%20Technical%20Review%20Panel%20Public%20Report.pdf>.
- . 2019. "U.S. Department of Energy Awards \$3.5 Million to New Gas Reactor Design." EnergyGov. December 5, 2019. <https://www.energy.gov/ne/articles/us-department-energy-awards-35-million-new-gas-reactor-design>.
- . 2020. "Energy Department's Advanced Reactor Demonstration Program Awards \$30 Million in Initial Funding for Risk Reduction Projects." Energy.gov. December 16, 2020. <https://www.energy.gov/ne/articles/energy-department-s-advanced-reactor-demonstration-program-awards-30-million-initial>.
- dos Santos Guimaraes, Leonam, and Jose Augusto Perrotta. 2020. "Brazil Looks to HALEU - Nuclear Engineering International." *Nuclear Engineering International*, June 24, 2020. <https://www.neimagazine.com/features/featurebrazil-looks-to-haleu-7993503>.

- Dubberly, A.E., T. Wu, and S. Kubo. 2003. "S-PRISM High Burnup Metal-Fuel Core Design: For Session 3: Future Deployment Programs And Issues." In *Proceedings of ICAPP '03*. Cordoba, Spain: GE Nuclear Energy.
- Duchnowski, E., N. Brown, Lance Snead, and J. Trelewicz. 2019. "Preliminary Reactor Physics Analysis of Alternative Moderators for Modular High Temperature Gas-Cooled Reactors." *Transactions* 121 (1): 1448–49. <https://doi.org/10.13182/T31307>.
- Durst, Phillip Casey, David Beddingfield, Brian Boyer, Robert Bean, Michael Cohen, Michael Ehinger, David Hanks, David L. Moses, and Lee Refalo. 2009. "Nuclear Safeguards Considerations For The Pebble Bed Modular Reactor (PBMR)." INL/EXT-09-16782, 968683. Idaho Falls, ID: Idaho National Laboratory. <https://doi.org/10.2172/968683>.
- Dyck, Elisabeth. 2018. "Towards More Sustainable Nuclear Energy with Non-Electric Applications: Opportunities and Challenges." International Atomic Energy Agency. December 20, 2018. <https://www.iaea.org/newscenter/news/towards-more-sustainable-nuclear-energy-with-non-electric-applications-opportunities-and-challenges>.
- Dyke, Elmer. 2020. "HALEU: The Path Forward." Presented at the GAIN-EPRI-NEI High-Assay Low-Enriched Uranium (HALEU) Webinar, April 28. https://gain.inl.gov/HALEU_Webinar_Presentations/13-Dyke,Centrus-28Apr2020.pdf.
- Ebbinghaus, B. B., C. G. Bathke, D. S. Dalton, and J. P. Murphy. 2013. "The Application of Material Attractiveness in a Graded Approach to Nuclear Materials Security." LLNL-CONF-640708. Livermore, CA: Lawrence Livermore National Laboratory. <https://www.osti.gov/biblio/1097225>.
- EIA (Energy Information Administration). 2021. "Annual Energy Outlook 2021." Department of Energy (DOE). https://www.eia.gov/outlooks/aeo/pdf/AEO_Narrative_2021.pdf.
- Elysium Industries. n.d. "The Molten Chloride Salt Fast Reactor." Technology. Accessed March 4, 2021. <https://web.archive.org/web/20201124135007/www.elysiumindustries.com/technology>.
- EPRI (Electric Power Research Institute). 2020. "Uranium Oxycarbide Tristructural Isotropic Coated Particle Fuel Performance." Topical Report EPRI-AR-1(NP)-A. <https://www.epri.com/research/programs/065093/results/3002019978>.
- Flanagan, George, Tom Fanning, and Tanju Sofu. 2015. "Sodium-Cooled Fast Reactor (SFR) Technology and Safety Overview." Presented at the Department of Energy Seminar, Washington, DC, February 18. <https://gain.inl.gov/SiteAssets/Fast%20Reactors/SFR-NRCTechnologyandSafetyOverview18Feb15.pdf>.
- Fletcher, David. 2020. "URENCO Next Generation Fuels: Conversion and Enrichment Options." Presented at the GAIN-EPRI-NEI High-Assay Low-Enriched Uranium (HALEU) Webinar, April 28. https://gain.inl.gov/HALEU_Webinar_Presentations/11-Fletcher,URENCO-28Apr2020.pdf.
- Ford, Michael J., Ahmed Abdulla, M. Granger Morgan, and David G. Victor. 2017. "Expert Assessments of the State of U.S. Advanced Fission Innovation." *Energy Policy* 108 (September): 194–200. <https://doi.org/10.1016/j.enpol.2017.05.059>.
- Forsberg, Charles W. 2006. "Molten-Salt-Reactors Technology Gaps." In *Proceedings of ICAPP '06*. Reno, NV: Oak Ridge National Laboratory. <https://www.osti.gov/etdweb/biblio/21021093>.
- Forsberg, C.W., C.M. Hopper, J.L. Richter, and H.C. Vantine. 1998. "Definition of Weapons-Usable Uranium-233." ORNL/TM-13517. Oak Ridge, TN: Oak Ridge National Laboratory. <https://thoriumenergyalliance.com/wp-content/uploads/2020/02/weapons-usable-u-233-ORNL-TM-13517.pdf>.
- Freebairn, Bill. 2020. "GE Hitachi, Gates' Firm Pitch Reactor Designed to Share Grid with Renewables." S&P Global Market Intelligence. September 11, 2020. <https://www.spglobal.com/marketintelligence/en/news-insights/latest-news-headlines/ge-hitachi-gates-firm-pitch-reactor-designed-to-share-grid-with-renewables-60276239>.
- Freebairn, William. 2019. "X-Energy Continues with Plan for Advanced Fuel Fabrication Facility, Executive Says." *NuclearFuel* 44 (4): 4–5.
- Freis, Daniel, Abdel El Abjani, Dragan Coric, Ramil Nasyrow, Joseph Somers, Chunhe Tang, Rongzheng Liu, Bing Liu, and Malin Liu. 2020. "Burn-up Determination and Accident Testing of HTR-PM Fuel Elements Irradiated in the HFR Petten." *Nuclear Engineering and Design* 357 (February): 110414. <https://doi.org/10.1016/j.nucengdes.2019.110414>.
- GA (General Atomics). 2019. "Advanced Reactors." 2019. <https://web.archive.org/web/20190413081031/http://www.ga.com/advanced-reactors>.
- . 2020. "General Atomics and Framatome Collaborate to Develop a Fast Modular Reactor." 2020. <https://web.archive.org/web/20190413081031/http://www.ga.com/advanced-reactors>.
- Gallacher, Vince. 2020. "TerraPower HALEU." Presented at the GAIN-EPRI-NEI High-Assay Low-Enriched Uranium (HALEU) Webinar, April 29. https://gain.inl.gov/HALEU_Webinar_Presentations/30-Gallacher,TerraPower-29Apr2020.pdf.
- GAO (Government Accountability Office). n.d. "Key Issues: Nuclear Waste Disposal." Nuclear Waste Disposal. Accessed March 4, 2021. https://www.gao.gov/key_issues/nuclear_waste_disposal/issue_summary?from=topics.
- Gehin, Jess C., and Jeffrey J. Powers. 2016. "Liquid Fuel Molten Salt Reactors for Thorium Utilization." *Nuclear Technology* 194 (2): 152–61. <https://doi.org/10.13182/NT15-124>.
- General Electric-Hitachi. 2021. "PRISM." GE Hitachi Nuclear Energy. 2021. <https://nuclear.gepower.com/build-a-plant/products/nuclear-power-plants-overview/prism1>.
- Gilleland, John, Robert Petroski, and Kevan Weaver. 2016. "The Traveling Wave Reactor: Design and Development." *Engineering* 2 (1): 88–96. <https://doi.org/10.1016/J.ENG.2016.01.024>.
- Gitau, Ernest Travis Ngure. 2011. "Development and Evaluation of a Safeguards System Concept for a Pebble-Fueled High Temperature Gas-Cooled Reactor." Texas A&M University. <https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.913.6697&rep=rep1&type=pdf>.
- Gougar, H.D., R.A. Bari, T.K. Kim, T.E. Sowinski, and A. Worrall. 2015. "Assessment of the Technical Maturity of Generation IV Concepts for Test or Demonstration Reactor Applications." INL/EXT-15-36427. Idaho Falls, ID: Idaho National Laboratory. <https://inldigitallibrary.inl.gov/sites/sti/sti/6721146.pdf>.
- Grabaskas, David, Matthew Bucknor, James Jerden, Acacia J. Brunett, Matthew Denman, Andrew Clark, and Richard S. Denning. 2016. "Regulatory Technology Development Plan—Sodium Fast Reactor: Mechanistic Source Term—Trial Calculation." ANL-ART-49-Voll-Vol2. Argonne, IL: Argonne National Laboratory. <https://doi.org/https://doi.org/10.2172/1334189>.

- Greenspan, Ehud. 2016. "A Pebble-Bed Breed-and-Burn Reactor." NEUP Project 13-5144. DOE (Department of Energy). <https://doi.org/10.2172/1244996>.
- Gronlund, Lisbeth, David Lochbaum, and Edwin Lyman. 2007. "Nuclear Power in a Warming World." Cambridge, MA: Union of Concerned Scientists. <https://www.ucsusa.org/resources/nuclear-power-warming-world#ucs-report-downloads>.
- Hayes, Steven L. 2017. "Metallic Fuels for Fast Reactors." Presented at the Generation IV International Forum, August 22. https://www.gen-4.org/gif/upload/docs/application/pdf/2017-08/geniv_webinar_metallic_fuels_for_fast_reactors_final_shayes.pdf.
- Heidet, F. 2019. "Current Status of the VTR Core Design." Presented at the European Research Reactor Conference, Swemeh, Jordan, March 24.
- Hejzlar, Pavel, Robert Petroski, Jesse Cheatham, Nick Touran, Michael Cohen, Bao Truong, Ryan Latta, et al. 2013. "Terrapower, LLC Traveling Wave Reactor Development Program Overview." *Nuclear Engineering and Technology* 45 (6): 731–44. <https://doi.org/10.5516/NET.02.2013.520>.
- Hoffman, E., and T. Fei. 2019. "Fast Reactor Design for Using Low-Enriched Uranium Startup and Transition." Presented at the GLOBAL International Nuclear Fuel Cycle Conference, Seattle, Washington, September 22.
- Hoffman, E.A., W.S. Yang, and R.N. Hill. 2008. "Preliminary Core Design Studies for the Advanced Burner Reactor over a Wide Range of Conversion Ratios." ANL-AFCI-177, 973480. Argonne, IL: Argonne National Laboratory. <https://doi.org/10.2172/973480>.
- Holcomb, D.E., W.P. Poore, and G.F. Flanagan. 2020. "MSR Fuel Salt Qualification Methodology." ORNL/TM-2020/157. Oak Ridge, TN: Oak Ridge National Laboratory. <https://adamswebsearch2.nrc.gov/webSearch2/view?AccessionNumber=ML2019A257>.
- Hombourger, B., A. Pautz, J. Krepel, and K. Mikityuk. 2015. "Fuel Cycle Analysis of a Molten Salt Reactor for Breed-and-Burn Mode - 15524." In *Proceedings of ICAPP '15*. Nice, France. http://inis.iaea.org/Search/search.aspx?orig_q=RN:49004874.
- Hou, Jason, Staffan Qvist, and Ehud Greenspan. 2015. "3-D Fuel Shuffling for Reduced Peak Burnup and Increased Uranium Utilization of Breed-and-Burn Reactors." In *Proceedings of ICAPP '15*. Nice, France. https://www.researchgate.net/profile/Jason-Hou-2/publication/277331547_3-D_Fuel_Shuffling_for_Reduced_Peak_Burnup_and_Increased_Uranium_Utilization_of_Breed-and-Burn_Reactors/links/5568b00a08aeccd77739ce17/3-D-Fuel-Shuffling-for-Reduced-Peak-Burnup-and-Increased-Uranium-Utilization-of-Breed-and-Burn-Reactors.pdf.
- Humrickhouse, Paul W. 2011. "HTGR Dust Safety Issues and Needs for Research and Development." INL/EXT-11-21097, 1023483. Idaho Falls, ID: Idaho National Laboratory. <https://doi.org/10.2172/1023483>.
- Hunn, John D., Grant W. Helmreich, Darren J. Skitt, and John A. Dyer. 2017. "Acceptance Test Data for BWXT Coated Particle Batch 93164A." ORNL/TM-2017/035. Oak Ridge, TN: Oak Ridge National Laboratory. <https://info.ornl.gov/sites/publications/Files/Pub72613.pdf>.
- Hunn, John, Bob Morris, Chuck Baldwin, Tyler Gerczak, Fred Montgomery, Darren Skitt, and John Dyer. 2017. "AGR-2 PIE and Safety Testing." Idaho Falls, ID, July 18. https://art.inl.gov/trisofuels/AGR%20TRISO%20Fuels%20Program%20Review%20July%201819%202017/06_Hunn_AGR-2_PIE_at_ORNL.pdf.
- Huotari, John. 2017. "A Great Technical Achievement, Molten Salt Reactor Could Be Entombed." Oak Ridge Today. November 26, 2017. <https://oakridgetoday.com/2017/11/26/great-technical-achievement-molten-salt-reactor-entombed>.
- IAEA (International Atomic Energy Agency). n.d. "Molten Salt Reactors." Text. n.d. <https://www.iaea.org/topics/molten-salt-reactors>.
- . 2011. "Energy, Electricity and Nuclear Power Estimates for the Period up to 2050." IAEA-RDS-1/31. Vienna, Austria: IAEA (International Atomic Energy Agency). https://www-pub.iaea.org/MTCD/Publications/PDF/RDSL_31.pdf.
- . 2014. "International Safeguards in the Design of Nuclear Reactors." No. NP-T-2.9. IAEA Nuclear Energy Series. Vienna, Austria. https://www-pub.iaea.org/MTCD/Publications/PDF/Pub1669_web.pdf.
- . 2016. "Status Report – IMSR-400." <https://aris.iaea.org/PDF/IMSR400.pdf>.
- . 2017. *Benchmark Analysis of EBR-II Shutdown Heat Removal Tests*. IAEA TECDOC Series 1819. Vienna, Austria: International Atomic Energy Agency. <https://public.ebookcentral.proquest.com/choice/publicfullrecord.aspx?p=5059780>.
- . 2018. "Development and Implementation Support Programme for Nuclear Verification 2018–2019." STR-386. <https://www.iaea.org/sites/default/files/18/09/sg-str-386-development-support-programme.pdf>.
- . 2020. "Energy, Electricity and Nuclear Power Estimates for the Period up to 2050." IAEA-RDS-1/40. Vienna, Austria: IAEA (International Atomic Energy Agency). <https://linkinghub.elsevier.com/retrieve/pii/0140670195951326>.
- INL (Idaho National Laboratory). n.d. "Frequently Asked Questions | Versatile Test Reactor." Idaho National Laboratory. n.d. <https://inl.gov/trending-topic/versatile-test-reactor/frequently-asked-questions>.
- . 2005. "Generation IV Nuclear Energy Systems Ten-Year Program Plan, Fiscal Year 2005." Volume I. Idaho Falls, ID: Idaho National Laboratory. <http://large.stanford.edu/courses/2015/ph240/morgan1/docs/gen-iv-10-yr-program-plan.pdf>.
- . 2012. "Assessment of High Temperature Gas-Cooled Reactor (HTGR) Capital and Operating Costs." TEV-1196. Idaho Falls, ID: Idaho National Laboratory. <https://art.inl.gov/NGNP/INL%20Documents/Year%202012/Assessment%20of%20High%20Temperature%20Gas-Cooled%20Reactor%20-%20HTGR%20-%20Capital%20and%20Operating%20Costs.pdf>.
- . 2020. "Materials and Fuels Complex FY-20 - FY-24 Five-Year Investment Strategy." INL/EXT-20-57248. Idaho Falls, ID: Idaho National Laboratory. https://inldigitallibrary.inl.gov/sites/sti/sti/Sort_23134.pdf.
- IRSN (Institut de Radioprotection et de Sûreté Nucléaire). 2015. "Review of Generation IV Nuclear Energy Systems." Fontenay-aux-Roses, France: IRSN (Institut de Radioprotection et de Sûreté Nucléaire). https://www.irsn.fr/EN/newsroom/News/Documents/IRSN_Report-GenIV_04-2015.pdf.
- Jian, Fu, Jiang Yingxue, Cheng Hongyong, and Cheng Wei. 2014. "Overview and Progress of High Temperature Reactor Pebble-bed Module Demonstration Project (HTR-PM)." In *Proceedings of the HTR 2014*. Weihai, China. <https://nucleus.iaea.org/sites/htrg-kb/HTR2014/Paper%20list/Track1/HTR2014-11125.pdf>.

- Jordan, Brett W., Roderick Eggert, Brent Dixon, and Brett Carlsen. 2014. "Thorium: Does Crustal Abundance Lead to Economic Availability?" Working Paper 2014-07. Division of Economics and Business Working Paper Series. Golden, CO: Colorado School of Mines. <http://econbus-papers.mines.edu/working-papers/wp201407.pdf>.
- Jubin, Robert T., Patricia D. Paviet, and Bresee, James C. 2016. "Xenon Recovery from Molybdenum-99 Production." *Procedia Chemistry* 21: 74–81. <https://www.sciencedirect.com/science/article/pii/S1876619616300535>.
- Kadak, Andrew C. 2016. "The Status of the US High-Temperature Gas Reactors." *Engineering 2* (1): 119–23. <https://doi.org/10.1016/J.ENG.2016.01.026>.
- Kang, Jungmin, and Frank von Hippel. 2005. "Limited Proliferation-Resistance Benefits from Recycling Unseparated Transuranics and Lanthanides from Light-Water Reactor Spent Fuel." *Science & Global Security* 13 (3): 169–181. http://scienceandglobalsecurity.org/archive/2005/09/limited_proliferation-resistan.html.
- Khatib-Rahbar, Mohsen, Alfred Krall, Zhe Yuan, and Michael Zavisca. 2020. "Review of Accident Tolerant Fuel Concepts with Implications to Severe Accident Progression and Radiological Releases." ERI/NRC 20-209. Rockville, Maryland: Energy Research, Inc. <https://www.nrc.gov/docs/ML2028/ML20287A477.pdf>.
- Kim, T. K., and T. A. Taiwo. 2010. "Fuel Cycle Analysis of Once-through Nuclear Systems." ANL-FCRD-308, 989083. Argonne, IL: Argonne National Laboratory. <https://doi.org/10.2172/989083>.
- King, T. L., R. R. Landry, E. D. Throm, and J. N. Wilson. 1991. "Pre-application Safety Evaluation Report for the Sodium Advanced Fast Reactor (SAFR) Liquid-Metal Reactor." NUREG-1369. Washington, DC: Nuclear Regulatory Commission (NRC). <https://www.nrc.gov/docs/ML0634/ML063410547.pdf>.
- Knol, S., F. Roelofs, N. Kothz, M. Laurie, M. Buckthorpe, W. Scheuermann, and P. Ho-Hune. 2015. "Final Report Summary—ARCHER (Advanced High-Temperature Reactors for Cogeneration of Heat and Electricity R&D)." EC (European Commission). <https://cordis.europa.eu/project/id/269892/reporting>.
- Kovacic, Donald N., Andrew Worrall, Louise G. Worrall, George F. Flanagan, David Eugene Holcomb, Robert Bari, Lap Cheng, David Farley, and Matthew Sternat. 2018. "Safeguards Challenges for Molten Salt Reactors." Oak Ridge, TN: Oak Ridge National Laboratory. <https://www.osti.gov/servlets/purl/1474868>.
- Kuperman, Alan J., David Sokolow, and Edwin S. Lyman. 2014. "Can the IAEA Safeguard Fuel-Cycle Facilities?" Working Paper # 2. Austin, TX: Nuclear Proliferation Prevention Project, LBJ School of Public Affairs, University of Texas at Austin. <http://sites.utexas.edu/nppp/files/2014/03/NPPP-working-paper-2-2014-Mar-18.pdf>.
- LeBlanc, David. 2010. "Molten Salt Reactors: A New Beginning for an Old Idea." *Nuclear Engineering and Design* 240 (6): 1644–56. <https://doi.org/10.1016/j.nucengdes.2009.12.033>.
- . 2016. "An Overview of the Integral Molten Salt Reactor." Presentation presented at the Oak Ridge National Laboratory Molten Salt Reactor Workshop, Oak Ridge, TN, October. <https://public.ornl.gov/conferences/MSR2016/docs/Presentations/MSR2016-day2-03-TEI-LEBLANC-ORNL-An-Overview-of-the-IMSR.pdf>.
- . 2021. "Leading the Way to a Bright Energy Future: Overview of the IMSR." Presented at the National Academy of Sciences Committee on Merits and Viability of Different Nuclear Fuel Cycles and Technology Options and the Waste Aspects of Advanced Nuclear Reactors, January 11. <https://www.nationalacademies.org/event/01-11-2021/docs/D4E2F679B3A899D11484719C1A812CE62308A2BD3CFF>.
- Lewis, E. E. 1977. *Nuclear Power Reactor Safety*. New York: John Wiley and Sons, Inc. https://inis.iaea.org/search/search.aspx?orig_q=RN:9379275.
- Li, Junfeng, Tao Ma, and Jianlong Wang. 2014. "Strategy Study on Treatment and Disposal of the Radioactive Graphite Waste of HTGR Spent Fuel." In *Proceedings of the HTR 2014*. Weihai, China. https://inis.iaea.org/collection/NCLCollectionStore/_Public/48/076/48076815.pdf.
- Lyman, Edwin. 2001. "The Pebble-Bed Modular Reactor (PBMR): Safety Issues." *Physics and Society* 30 (4): 16–19. <https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.491.815&rep=rep1&type=pdf>.
- . 2013. "Small Isn't Always Beautiful: Safety, Security, and Cost Concerns about Small Modular Reactors." Cambridge, MA: Union of Concerned Scientists. https://www.ucsusa.org/resources/small-modular-reactors?_ga=2.165740926.1323814253.1615749426-1169742399.1612975494.
- . 2017. "The Pyroprocessing Files." *All Things Nuclear* (blog). August 12, 2017. <https://allthingsnuclear.org/elyman/the-pyroprocessing-files>.
- . 2018a. "UCS Technical Rebuttal to the Idaho National Laboratory's Opinions on the Versatile (Fast) Test Reactor." Cambridge, MA: Union of Concerned Scientists. <https://s3.amazonaws.com/ucs-documents/global-security/Lyman-Response-INL-Justification.pdf>.
- . 2018b. "Advanced Reactors: Proliferation and Terrorism Concerns." APS Forum on Physics & Society (blog). October 10, 2018. <https://physicsandsocietyforum.wordpress.com/2018/10/10/advanced-reactors-proliferation-and-terrorism-concerns>.
- . 2019. "Proliferation and Terrorism Issues Related to Noble Gas Management at Molten Salt Reactors." In *Proceedings of Global/Top Fuel 2019*. Seattle, Washington.
- Lynas, Mark. 2011. "Good Reasons Not to Waste Nuclear 'Waste.'" *Mark Lynas: Environmental News and Comment* (blog). July 22, 2011. <https://www.marklynas.org/2011/07/good-reasons-not-to-waste-nuclear-waste>.
- Machiels, A., S. Massara, and C. Garzenne. 2009. "Dynamic Analysis of a Deployment Scenario of Fast Burner Reactors in the U.S. Nuclear Fleet." In *Proceedings of the Global 2009 Conference*. Paris, France.
- Maden, Claire. 2009. "No Limits for Yucca Mountain?" *Nuclear Engineering International*, February 20, 2009. <https://www.neimagazine.com/features/featureno-limits-for-yucca-mountain-721>.
- Mark, Carson. 1984. Conversion of Research and Test Reactors to Low-Enriched Uranium (LEU) Fuel. U.S. Government Printing Office. https://books.google.com/books?id=aKCvfvYzBgC&pg=PA295&lpq=PA295&dq=RERTR+hearing+1984&source=bl&ots=hWOP9twyD5&sig=ACfU3U1ve3WGUHnLah7_GqJipyWVZdK_oQ&hl=en&sa=X&ved=2ahUKewiJnoOtyqfvAhWWAZOJHeFmBds4ChDoATAHegQIBxAD#v=onepage&q=RERTR%20hearing%201984&f=false.

- Marshall, Douglas W. 2019. "AGR-5/6/7 Fuel Fabrication Report." INL/EXT-19-53720. Idaho Falls, ID: INL (Idaho National Laboratory). https://indigitalibrary.inl.gov/sites/sti/sti/Sort_14786.pdf.
- Martin, Michael, Manuele Aufiero, Ehud Greenspan, and Massimiliano Fratoni. 2017. "Feasibility of a Breed-and-Burn Molten Salt Reactor." In *Transactions of the American Nuclear Society*. Vol. 116. San Francisco, California. https://www.researchgate.net/publication/317489914_Feasibility_of_a_Breed-and-Burn_Molten_Salt_Reactor.
- Martin, W.R., J.C. Lee, E.W. Larsen, and M.C. Edlund. 1991. "The Slightly-Enriched Spectral Shift Control Reactor." DOE/ER/12809-4. Ann Arbor, MI: University of Michigan. <https://doi.org/10.2172/6076324>.
- Mathieu, L., D. Heuer, R. Brissot, C. Garzenne, C. Le Brun, D. Lecarpentier, E. Liatard, et al. 2006. "The Thorium Molten Salt Reactor: Moving on from the MSBR." *Progress in Nuclear Energy* 48 (7): 664–79. <https://doi.org/10.1016/j.pnucene.2006.07.005>.
- Mathieu, L., D. Heuer, E. Merle-Lucotte, R. Brissot, C. Le Brun, E. Liatard, J.-M. Loiseaux, O. Méplan, A. Nuttin, and D. Lecarpentier. 2009. "Possible Configurations for the Thorium Molten Salt Reactor and Advantages of the Fast Nonmoderated Version." *Nuclear Science and Engineering* 161 (1): 78–89. <https://doi.org/10.13182/NSE07-49>.
- McDowell, Bruce K., James R. Nickolaus, Mark R. Mitchell, Gary L. Swearingen, and Ray Pugh. 2011. "High Temperature Gas Reactors: Assessment of Applicable Codes and Standards." PNNL-20869, 1031438. Richland, WA: Pacific Northwest National Laboratory. <https://doi.org/10.2172/1031438>.
- McDuffee, Joel. 2017. "MSR/FHR Gap Analysis, Path Forward, and Q&A." Presented at the 2017 GAIN Fuel Safety Research Workshop, Idaho Falls, Idaho, May 3. <https://gain.inl.gov/SiteAssets/Fuel%20Safety%20Presentations/14%20McDuffee-MSR-FHR%20Gap%20Analysis.pdf>.
- McFarlane, J., N. D. Bull Ezell, G.D. DelCul, D. E. Holcomb, K. Myhre, A. Lines, S. Bryan, H. Felmy, and B. J. Riley. 2019. "Fission Product Volatility and Off-Gas Systems for Molten Salt Reactors." ORNL/TM-2019/1266; PNNL-28974. Oak Ridge, TN: Oak Ridge National Laboratory. <https://doi.org/10.2172/1560441>.
- McMillan, Bill. 2019. "Molten Salt Reactor Experiment Project Initiatives." Slides presented at the Oak Ridge Site Specific Advisory Board Monthly Meeting, Oak Ridge, TN, November 13. <https://www.energy.gov/sites/prod/files/2019/11/f68/ORSS-AB%20MSRE%20Presentation%20-%20Nov%202019.pdf>.
- Messick, N.C., P.R. Betten, W.F. Booty, L.J. Christensen, R.M. Fryer, D. Mohr, H.P. Planchon, and W.H. Radtke. 1987. "Modification of EBR-II Plant to Conduct Loss-of-Flow-without-Scram Tests." *Nuclear Engineering and Design* 101 (1): 13–23. [https://doi.org/10.1016/0029-5493\(87\)90146-4](https://doi.org/10.1016/0029-5493(87)90146-4).
- Mickum, G. Spencer, Robert D. McElroy, and Nolan E. Hertel. 2014. "Safeguards Motivations and Limitations on Pyroprocessing Activities." Presented at the Annual Meeting of the Institute of Nuclear Materials Management, Atlanta, GA, July 20.
- Milko, John, Todd Allen, and Ryan Fitzpatrick. 2018. "Keeping Up with the Advanced Nuclear Industry." Third Way. February 8, 2018. <https://www.thirdway.org/graphic/keeping-up-with-the-advanced-nuclear-industry>.
- MIT (Massachusetts Institute of Technology). 2011. *The Future of Nuclear Fuel Cycle: An Interdisciplinary MIT Study*. Cambridge, MA: MIT. http://large.stanford.edu/courses/2012/ph241/schultz2/docs/The_Nuclear_Fuel_Cycle-all.pdf.
- Moormann, R. 2009. "AVR Prototype Pebble Bed Reactor: A Safety Re-Evaluation of Its Operation and Consequences for Future Reactors." *Kerntechnik* 74: 8–21. <https://doi.org/10.3139/124.110001>.
- Moormann, Rainer, R. Scott Kemp, and Ju Li. 2018. "Caution Is Needed in Operating and Managing the Waste of New Pebble-Bed Nuclear Reactors." *Joule* 2 (10): 1911–14. <https://doi.org/10.1016/j.joule.2018.07.024>.
- Morgan, M. Granger, Ahmed Abdulla, Michael J. Ford, and Michael Rath. 2018. "US Nuclear Power: The Vanishing Low-Carbon Wedge." *Proceedings of the National Academy of Sciences* 115 (28): 7184–89. <https://doi.org/10.1073/pnas.1804655115>.
- Morris, R.N., D.A. Petti, D.A. Powers, and B.E. Boyack. 2004. "TRISO-Coated Particle Fuel Phenomenon Identification and Ranking Tables (PIRTs) for Fission Product Transport Due to Manufacturing, Operations, and Accidents: Main Report." NUREG/CR-6844, Vol. 1. Nuclear Regulatory Commission. <https://www.nrc.gov/docs/ML0423/ML042300668.pdf>.
- Mulder, E.J., and W.A. Boyes. 2020. "Neutronics Characteristics of a 165 MWth Xe-100 Reactor." *Nuclear Engineering and Design* 357: 110415. <https://doi.org/10.1016/j.nucengdes.2019.110415>.
- Nalley, Stephen, and Angelina LaRose. 2021. "Annual Energy Outlook 2021." Presented at the Annual Energy Outlook, Washington, DC, February 3. https://www.eia.gov/pressroom/presentations/AEO2021_Release_Presentation.pdf.
- Napier, Bruce A. 2020. "Non-LWR Environmental Data." PNNL-29367 Rev. 1. Richland, WA: Pacific Northwest National Laboratory. <https://www.nrc.gov/docs/ML2007/ML20076G075.pdf>.
- National Research Council. 1996. *Nuclear Wastes: Technologies for Separations and Transmutation*. Washington, DC: The National Academies Press. <https://doi.org/10.17226/4912>.
- NEA (Nuclear Energy Agency). 2004. "Non-Electricity Products of Nuclear Energy." Organisation for Economic Co-Operation And Development. <https://www.oecd-neo.org/ndd/reports/2004/non-electricity-products.pdf>.
- . 2012. *Actinide and Fission Product Partitioning and Transmutation*. Nuclear Science. Paris, France: OECD NEA. <https://doi.org/10.1787/9789264991743-en>.
- . 2015. "Perspectives on the Use of Thorium in the Nuclear Fuel Cycle: Extended Summary." NEA No. 7228. Issy-les-Moulineaux, France: Organisation for Economic Co-Operation and Development. https://www.oecd-neo.org/jcms/pl_14932.
- NEA (Nuclear Energy Agency), and IAEA (International Atomic Energy Agency). 2020. "Uranium Resources, Production and Demand." https://www.oecd-neo.org/jcms/115673_Media/red-book-2020?details=true.
- NEI (Nuclear Engineering International). 2020. "Georgia Power Orders Nuclear Fuel for Vogtle 4." *Nuclear Engineering International*, March 9, 2020. <https://www.neimagazine.com/news/newsgeorgia-power-orders-nuclear-fuel-for-vogtle-4-7813692>.
- NGNP (Next Generation Nuclear Plant) Industry Alliance Limited. 2021. "The High Temperature Gas-Cooled Reactor (HTGR)—Safe, Clean and Sustainable Energy for the Future." <https://web.archive.org/web/20190619231231/https://www.ngnpalliance.org/index.php/htgr>.

- Nicobar Group, and US Commercial Service. 2017. "China's Nuclear Industry 2017-2018: A Tightly Coiled Spring." https://static1.squarespace.com/static/54a897dbe4b0631d0c50cdbc/t/5a5d7fed71c10bc094350a36/1516077082229/2017-2018_China+Nuclear+Power+White+Paper_+Nicobar+Group.pdf.
- NRC (Nuclear Regulatory Commission). 1994. Preapplication Safety Evaluation Report for the Power Reactor Innovative Small Module (PRISM) Liquid-Metal Reactor. NUREG-1368. Washington, DC. www.nrc.gov/docs/ML0634/ML063410561.pdf.
- . 2012. "Assessment of White Paper Submittals on Fuel Qualification and Mechanistic Source Terms." Project 0748. Washington, DC. <https://www.nrc.gov/docs/ML1202/ML120240669.pdf>.
- . 2017. "A Regulatory Review Roadmap for Non-Light Water Reactors." ML17312B567. Washington, DC. <https://www.nrc.gov/docs/ML1731/ML17312B567.pdf>.
- . 2019. Proposed Rule - Emergency Preparedness for Small Modular Reactors and Other New Technologies (RIN 3150 AJ68; NRC-2015-0225). SRM-SECY-18-0102. <https://www.nrc.gov/docs/ML1935/ML19351C729.html>.
- . 2020a. "Final Safety Evaluation: Uranium Oxycarbide (UCO) Tristructural Isotropic (TRISO) Coated Particle Fuel Performance: Topical Report EPRI-AR-1(NP)." Washington, DC: NRC (Nuclear Regulatory Commission). www.nrc.gov/ML2021/ML20216A453.pdf.
- . 2020b. "U.S Nuclear Regulatory Commission Staff Observations on Terrestrial Energy's Integral Molten Salt Reactor Core-Unit Interfaces White Paper." <https://www.nrc.gov/docs/ML2019/ML20191A384.pdf>.
- Obayashi, Yuka, and Aaron Sheldrick. 2018. "Japan Pledges to Cut Plutonium Stocks amid Growing Concern from Neighbors." *Reuters*, July 31, 2018. <https://www.reuters.com/article/us-japan-nuclear-plutonium/japan-pledges-to-cut-plutonium-stocks-amid-growing-concern-from-neighbors-idUSKBN1KL0I4>.
- Ohno, Kaoru. 2020. "NRA Grants Permission to Make Changes to Reactor Installation at HTTR under New Regulatory Standards." Japan Atomic Industrial Forum. June 2020. <https://www.jaif.or.jp/en/nra-grants-permission-to-make-changes-to-reactor-installation-at-htr-under-new-regulatory-standards>.
- Oklo. 2020. "Part II: Final Safety Analysis Report." OkloPower-2020-PartII-NP, Rev. 0. Sunnyvale, CA. <https://www.nrc.gov/docs/ML2007/ML20075A003.pdf>.
- Paillère, Henri. 2013. "Past NEA Activity and Objectives of the Workshop." Slides presented at the Joint IAEA/NEA Expert Workshop on the Technical and Economic Assessment of Non-Electric Applications of Nuclear Energy, Paris, France, April 4. https://www.oecd-nea.org/hdd/workshops/nucogen/presentations/1_Paillere_Technical-Economic-Assessment.pdf.
- Palmer, Joe. 2019. "AGR-5/6/7 Experiment Status as of the End of Cycle 166A." INL/MIS-19-56555-Revision-0. Idaho Falls, ID: Idaho National Laboratory. https://inldigitalibrary.inl.gov/sites/sti/sti/Sort_20504.pdf.
- . 2020. "AGR-5/6/7 Irradiation Summary as of the End of Cycle 167A." Slides presented at the Gas-Cooled Reactor Program Annual Review, Idaho Falls, ID, July 14. https://art.inl.gov/Meetings/GCR%20Program%20Review%20July%202020/Presentations/Session%202/04_PALMER_AGR%205-6-8%20Irradiation%20Summary.pdf.
- Patel, Sonal. 2019. "France Scraps Fast Nuclear Reactor Demonstration." *POWER Magazine*, November 1, 2019. <https://www.powermag.com/france-scraps-fast-nuclear-reactor-demonstration>.
- Paviet, Patricia. 2017. "Department of Energy's Efforts on the R&D Related to the Back End of the Nuclear Fuel Cycle." Presented at the American Nuclear Society Winter Meeting, Washington, DC, November 1.
- Pellum, Travis. 2019. "L.F.T.R. the New Age of Nuclear." Instructables. January 17, 2019. <https://www.instructables.com/LFTR-the-New-age-of-Nuclear>.
- Petersen, Eirik Eide. 2016. "Coupled Multi-Physics Simulations of the Molten Salt Fast Reactor Using Coarse-Mesh Thermal-Hydraulics and Spatial Neutronics." Master's Thesis, Saint-Aubin, France: Université Paris-Saclay. <https://www.psi.ch/sites/default/files/import/fast/PublicationsEN/FB-DOC-16-016.pdf>.
- Petti, D. A., J. Buongiorno, J. T. Maki, and G. K. Miller. 2002. "Key Differences in the Fabrication of US and German TRISO-Coated Particle Fuel, and Their Implications on Fuel Performance." In *Proceedings of HTR-2002*. Petten, Netherlands. https://inis.iaea.org/collection/NCLCollectionStore/_Public/33/033/33033032.pdf.
- Petti, D., R. Hill, J. Gehin, H. Gougar, G. Strydom, F. Heidet, J. Kinsey, et al. 2017. "Advanced Demonstration and Test Reactor Options Study." INL/EXT-16-37867, Revision 3. Idaho Falls, ID: Idaho National Laboratory. https://art.inl.gov/ART%20Document%20Library/Advanced%20Demonstration%20and%20Test%20Reactor%20Options%20Study/ADTR_Options_Study_Rev3.pdf.
- Petti, David, Tim Abram, Richard Hobbins, and Jim Kendall. 2010. "Updated NGNP Fuel Acquisition Strategy." INL/EXT-07-12441, Rev. 2. Idaho Falls, ID: Idaho National Laboratory. <https://doi.org/10.2172/1009165>.
- Petti, David, Collin Blaise, and Douglas Marshall. 2017. "A Summary of the Results from the DOE Advanced Gas Reactor (AGR) Fuel Development and Qualification Program." INL/EXT-16-40784, 1369357. Idaho Falls, ID: Idaho National Laboratory. <https://doi.org/10.2172/1369357>.
- Pham, Binh T., and Dawn M. Scates. 2019. "AGR-5/6/7 Release-to-Birth Ratio Data Analysis for Cycles 162B, 163A, 164A, and 164B." INL/EXT-19-54457, Revision 0. Idaho Falls, ID: INL (Idaho National Laboratory). https://inldigitalibrary.inl.gov/sites/sti/sti/Sort_17283.pdf.
- Pham, Binh T., James W. Sterbentz, Grant L. Hawkes, Dawn M. Scates, and Joe. Palmer. 2020. "AGR-5/6/7 Experiment Monitoring and Simulation Progress." INL/EXT-19-55429, Revision 1. Idaho Falls, ID: Idaho National Laboratory. <https://doi.org/10.2172/1599405>.
- Piet, Steven J., Samuel E. Bays, and Nick Soelberg. 2010. "HTGR Technology Family Assessment for a Range of Fuel Cycle Missions." INL/EXT-10-19704. Idaho Falls, ID: Idaho National Laboratory. <https://inldigitalibrary.inl.gov/sites/sti/sti/4731793.pdf>.
- Piet, Steven J., Brent W. Dixon, Jacob J. Jacobson, Gretchen E. Matthern, and David E. Shropshire. 2011. "Dynamic Simulations of Advanced Fuel Cycles." *Nuclear Technology* 173 (3): 227–38. <https://doi.org/10.13182/NT11-A11658>.

- Piet, Steven J., Edward A. Hoffman, and Samuel E. Bays. 2010. "Implications of Fast Reactor Transuranic Conversion Ratio." In *Actinide and Fission Product Proceedings of the Partitioning and Transmutation Eleventh Information Exchange Meeting*. San Francisco, CA. <https://doi.org/10.1787/9789264991743-en>.
- Piet, Steven J., Edward A. Hoffman, Samuel E. Bays, Gretchen E. Matthern, Jacob J. Jacobson, Ryan Clement, and David W. Gerts. 2013. "Fuel Cycle System Analysis Implications of Sodium-Cooled Metal-Fueled Fast Reactor Transuranic Conversion Ratio." *Nuclear Technology* 181 (3): 427–58. <https://doi.org/10.13182/NT13-A15802>.
- Piet, Steven J., Nick R. Soelberg, Robert S. Cherry, Layne F. Pincock, Eric L. Shaber, Melissa C. Teague, Gregory M. Teske, et al. 2010. "Analyzing Losses: Transuramics Into Waste and Fission Products Into Recycled Fuel." Preprint INL/CON-10-20136. Idaho Falls, ID: Idaho National Laboratory. <https://inldigitallibrary.inl.gov/sites/sti/sti/4654911.pdf>.
- Planchon, H. P., G. H. Golden, J. I. Sackett, D. Mohr, L. K. Chang, E. E. Feldman, and P. R. Betten. 1988. "Results and Implications of the Experimental Breeder Reactor II Inherent Safety Demonstration Tests." *Nuclear Science and Engineering* 100 (4): 549–57. <https://doi.org/10.13182/NSE100-549>.
- Planchon, H.P., J.I. Sackett, G.H. Golden, and R.H. Sevy. 1987. "Implications of the EBR-II Inherent Safety Demonstration Test." *Nuclear Engineering and Design* 101 (1): 75–90. [https://doi.org/10.1016/0029-5493\(87\)90152-X](https://doi.org/10.1016/0029-5493(87)90152-X).
- Price, Terry J., Ondrej Chvala, and Zack Taylor. 2020. "Xenon in Molten Salt Reactors: The Effects of Solubility, Circulating Particulate, Ionization, and the Sensitivity of the Circulating Void Fraction." *Nuclear Engineering and Technology* 52 (6): 1131–36. <https://doi.org/10.1016/j.net.2019.11.026>.
- Qualls, Lou, and David Holcomb. n.d. "Molten Salt Reactor Instrumentation." Accessed March 11, 2021. <https://www.energy.gov/sites/prod/files/2019/12/f69/ne-05-2019-Molten-Salt-Reactors-Overview-Qualls.pdf>.
- Qvist, Staffan, and Ehud Greenspan. 2012. "Inherent Safety of Minimum-Burnup Breed & Burn Reactors." In *Proceedings of the 2012 International Congress on Advances in Nuclear Power Plants*. Chicago, IL: American Nuclear Society. https://www.researchgate.net/publication/261172427_Inherent_safety_of_minimum-burnup_breed_burn_reactors.
- RAHP (Research Association of High Temperature Gas Cooled Reactor Plant). 2019. "High Temperature Gas Cooled Reactor (HTGR) Developments in the World - Present Status and Future Plans." Tokyo, Japan: Research Association of High Temperature Gas Cooled Reactor Plant. http://www.iae.or.jp/htgr/pdf/05_newsletter/05_2/05_2_rahp_No18_201903_en.pdf.
- Redmond, Everett II. 2020. "HALEU Needs For Commercial Reactors." Slides presented at the GAIN-EPRI-NEI High-Assay Low-Enriched Uranium (HALEU) Webinar, Webinar, April 28. https://gain.inl.gov/HALEU_Webinar_Presentations/06-Redmond,Industry_HALEU_Projected_Needs,Rev5-28Apr2020.pdf.
- Riley, Brian J., Joanna McFarlane, Guillermo D. DelCul, John D. Vienna, Cristian I. Contescu, and Charles W. Forsberg. 2019. "Molten Salt Reactor Waste and Effluent Management Strategies: A Review." *Nuclear Engineering and Design* 345: 94–109. <https://doi.org/10.1016/j.nucengdes.2019.02.002>.
- Robertson, Sean, Leslie Dewan, Mark Massie, Eva E. Davidson, Benjamin R. Betzler, Andrew Worrall, and Jeffrey J. Powers. 2017. "Assessment of the Neutronic and Fuel Cycle Performance of the Transatomic Power Molten Salt Reactor Design." Technical Report ORNL/TM-2017/475, CRADA/NFE-16-06345. Oak Ridge, TN: Oak Ridge National Laboratory. <https://doi.org/10.2172/1410921>.
- Rodrigues, Davide, Gabriela Durán-Klie, and Sylvie Delpech. 2015. "Pyrochemical Reprocessing of Molten Salt Fast Reactor Fuel: Focus on the Reductive Extraction Step." *Nukleonika* 60 (4): 907–14. <https://doi.org/10.1515/nuka-2015-0153>.
- Rogelj, Joeri, Drew Shindell, Kejun Jiang, Solomon Fifita, Piers Forster, Veronika Ginzburg, Collins Handa, et al. 2018. "Mitigation Pathways Compatible with 1.5°C in the Context of Sustainable Development." In *Global Warming of 1.5 °C*. IPCC. <https://www.ipcc.ch/sr15/chapter/chapter-2>.
- Ruggieri, J-M., L. Ren, J-P. Glatz, I. Ashurko, H. Hayafune, Y. Kim, and R. Hill. 2017. "Sodium-Cooled Fast Reactor (SFR) System Safety Assessment." Generation IV International Forum. https://www.gen-4.org/gif/upload/docs/application/pdf/2017-11/gif-sfr-safetyassessment-20170427_final.pdf.
- Schneider, Mycle, Antony Frogatt, Julie Hazemann, Ali Ahmad, Tadahiro Katsuta, M V Ramana, Ben Wealer, Agnès Stienne, Friedhelm Meinass, and Jungmin Kang. 2020. "The World Nuclear Industry Status Report 2020." <https://www.worldnuclearreport.org/-World-Nuclear-Industry-Status-Report-2020-.html>.
- Scott, Ian R. 2017. "The Stable Salt Reactor: Transforming the Promise of the Molten Salt Fuel Concept into a Viable Technology." *Modern Power Systems*, February 23, 2017. <https://modernpowersystems.com/features/featurethe-stable-salt-reactor-transforming-the-promise-of-the-molten-salt-fuel-concept-into-a-viable-technology-5748063>.
- Seaborg Technologies. n.d. "Seaborg: Making Nuclear Sustainable." White Paper. <https://drive.google.com/file/d/0Bx2cC35KJTwSDU1yUDNjd3NYOXREOFRuWDAzU2hpUjlfUHRn/view>.
- . 2015. "Seaborg Wasteburner: Molten Salt Reactor." White Paper SEAB-WP-2015-001. <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.712.5069&rep=rep1&type=pdf>.
- Serp, Jérôme, Michel Allibert, Ondřej Beneš, Sylvie Delpech, Olga Feynberg, Véronique Ghetta, Daniel Heuer, et al. 2014. "The Molten Salt Reactor (MSR) in Generation IV: Overview and Perspectives." *Progress in Nuclear Energy* 77: 308–19. <https://doi.org/10.1016/j.pnucene.2014.02.014>.
- Shafer, Deva. 2018. "Design and Melting Behavior of the MSFR Freeze Plug." Master's thesis, Delft, Netherlands: Delft University of Technology. <http://resolver.tudelft.nl/uuid:a2032f51-6948-40f5-ba35-5b40309d1fa2>.
- Shropshire, David E., and J. Stephen Herring. 2004. "Fuel-Cycle and Nuclear Material Disposition Issues Associated with High-Temperature Gas Reactors." In *Proceedings of the American Nuclear Energy Symposium (ANES 2004)*. Miami Beach, FL. <https://www.osti.gov/servlets/purl/839371>.
- Soltoff, Ben. 2020. "Advanced Nuclear: A Climate-Tech Comeback Story." *GreenBiz*, November 3, 2020. <https://www.greenbiz.com/article/advanced-nuclear-climate-tech-comeback-story>.

- Sterbentz, J. W., P. D. Bayless, L. Nelson, H. D. Gougar, J. Kinsey, and G. Strydom. 2016. "High Temperature Gas-Cooled Test Reactor Point Design: Summary Report." INL/EXT-16-37661. Idaho Falls, ID: Idaho National Laboratory. <https://inldigitallibrary.inl.gov/sites/sti/sti/6748415.pdf>.
- Steven J. Piet, Samuel E. Bays, and Nick Soelberg. 2010. "HTGR Technology Family Assessment for a Range of Fuel Cycle Missions." INL/EXT-10-19704, 1004243. Idaho National Laboratory. <https://doi.org/10.2172/1004243>.
- Stone, Robert. 2013. *Pandora's Promise*. Documentary. CNN. <http://transcripts.cnn.com/TRANSCRIPTS/1311/07/se.01.html>.
- Sumner, Tyler, and Thomas Fanning. 2020. "Versatile Test Reactor Pump Coastdown Assessment." In *Transactions of the American Nuclear Society*, 123:987–88.
- Sunseri, Matthew W. Letter report. 2020. "Safety Evaluation, Revision 2 for Topical Report, 'Uranium Oxycarbide (UCO) Tristructural Isotropic (TRISO) Coated Particle Fuel Performance: Topical Report EPRI-AR-1 (NP),' " August 4, 2020. <https://www.nrc.gov/docs/ML2021/ML20210M241.pdf>.
- Tang, Chunhe, Xiaoming Fu, Junguo Zhu, Tongxiang Liang, Konstantin N. Koshcheyev, Alexandre V. Kozlov, Oleg G. Karlov, Yu. G. Degaltsev, and Vladimir I. Vasiliev. 2006. "Fuel Irradiation of the First Batches Produced for the Chinese HTR-10." *Nuclear Engineering and Design* 236 (1): 107–13. <https://doi.org/10.1016/j.nucengdes.2005.11.001>.
- Tavron, Barak, and Eugene Shwageraus. 2016. "Pebble Bed Reactor Fuel Cycle Optimization Using Particle Swarm Algorithm." *Nuclear Engineering and Design* 307 (October): 96–105. <https://doi.org/10.1016/j.nucengdes.2016.06.033>.
- Temple, James. 2017. "Nuclear Energy Startup Transatomic Backtracks on Key Promises." *MIT Technology Review*, February 24, 2017. <https://www.technologyreview.com/2017/02/24/68882/nuclear-energy-startup-transatomic-backtracks-on-key-promises>.
- Tentner, A M, E Parma, T Wei, R Wigeland, Nuclear Engineering Division, SNL, and INL. 2010. "Severe Accident Approach—Final Report. Evaluation of Design Measures for Severe Accident Prevention and Consequence Mitigation." Argonne, IL: Argonne National Laboratory. <https://doi.org/10.2172/973483>.
- Terrani, Kurt. 2019. "Transformational Challenge Reactor Program." Presented at the Advanced Reactors Summit VI & Technology Trailblazers Showcase, La Jolla, CA, January 31.
- Terrestrial Energy. n.d. "IMSR®: Transformational Clean Energy Technology." Terrestrial Energy. n.d. <http://www.terrestrialenergy.com>.
- . 2020. "IMSR® Core-Unit Definition: Applicable Structures, Systems and Components." TEUSA Document #:200310. Charlotte, NC: Terrestrial Energy USA. <https://www.nrc.gov/docs/ML2009/ML20097B839.pdf>.
- . 2021. "TEUSA Response to NRC Feedback Questions and Observations Related to the TEUSA Principal Design Criteria White Paper." TEUSA Document#: R210106. Charlotte, NC: Terrestrial Energy USA. <https://www.nrc.gov/docs/ML2102/ML21025A033.pdf>.
- Thomas, Steve. 2008. "Safety Issues with the South African Pebble Bed Modular Reactor: When Were the Issues Apparent? A Briefing Paper." London, England: University of Greenwich. <https://core.ac.uk/download/pdf/67307.pdf>.
- ThorCon Power. 2017. "Uranium Requirements." Thorcon Power. 2017. <https://web.archive.org/web/20170428220138/http://thorconpower.com/library/documents/uranium>.
- . 2018. "ThorCon: The Do-Able Molten Salt Reactor." <https://thorconpower.com/docs/domsr20180119.pdf>.
- Tiberga, Marco, Devaja Shafer, Danny Lathouwers, Martin Rohde, and Jan Leen Kloosterman. 2019. "Preliminary Investigation on the Melting Behavior of a Freeze-Valve for the Molten Salt Fast Reactor." *Annals of Nuclear Energy* 132: 544–54. <https://doi.org/10.1016/j.anucene.2019.06.039>.
- TAP (Transatomic Power Corporation). 2016a. "Neutronics Overview." Cambridge, MA. <http://www.transatomicpower.com/wp-content/uploads/2015/04/Neutronics-White-Paper-v1.1.pdf>.
- TAP (Transatomic Power Corporation). 2016b. "Technical White Paper, V. 2.1." Cambridge, MA: Transatomic Power.
- U-Battery. 2019. "U-Battery: Local Modular Energy: A Future Energy Solution." Prospectus. Stoke Poges, England: UBattery. https://www.u-battery.com/cdn/uploads/supporting-files/U-Battery_Prospectus_May20.pdf.
- Van Tuyle, G.J., G.C. Slovik, B.C. Chan, A.L. Aronson, and R.J. Kennett. 1992. "Evaluations of 1990 PRISM Design Revisions." Upton, NY: Brookhaven National Laboratory. <https://doi.org/10.2172/10137887>.
- Vietti-Cook, Annette L. 2018. "Staff requirements—SECY-18-0096—Functional Containment Performance Criteria for Non-Light-Water-Reactors." Memorandum to Margaret M. Doane, NRC Executive Director for Operations, December 4. Washington, DC: Nuclear Regulatory Commission. <https://www.nrc.gov/docs/ML1833/ML18338A502.pdf>.
- Walton, Rod. 2018. "Centrus, X-Energy Deepen Partnership for Advanced Nuclear Fuel and Reactor Work." *Power Engineering*, November 30, 2018. <https://www.power-eng.com/nuclear/centrus-x-energy-deepen-partnership-for-advanced-nuclear-fuel-and-reactor-work/#gref>.
- Wang, S., M. Massone, A. Rineiski, and E. Merle-Lucotte. 2016. "Analytical Investigation of the Draining System for a Molten Salt Fast Reactor." In *Proceedings of NUTHOS-11*. Gyeongju, South Korea. http://samofar.eu/wp-content/uploads/2017/01/2016_Wang_NUTHOS-11-paper.pdf.
- Westphal, B. R., J. C. Price, K. C. Marsden, and M. N. Patterson. 2017. "Experience of Graphite Crucibles during the Pyrometallurgical Processing of Uranium." In *Proceedings of GLOBAL 2017*. Seoul, Korea: Idaho National Laboratory. <https://www.osti.gov/biblio/1411593>.
- Whiting, N., M. Hosoya, and J. Doo. 2006. "Development of a state-level integrated safeguards approach for Canada." In *Proceedings of the 47th Annual Meeting of the Institute of Nuclear Materials Management*, Nashville, TN, July 16-20, 2006.
- Wigeland, R., T. Taiwo, H. Ludewig, M. Todosow, W. Halsey, J. Gehin, R. Jubin, et al. 2014. "Nuclear Fuel Cycle Evaluation and Screening—Final Report." INL/EXT-14-31465. Washington, DC: Department of Energy. <https://fuelcycleevaluation.inl.gov/Shared%20Documents/ES%20Main%20Report.pdf>.
- Wigeland, R., T. Taiwo, M. Todosow, W. Halsey, and J. Gehin. 2009. "AFCI Options Study." INL/EXT-10-17639. Idaho Falls, ID: Idaho National Laboratory. <https://doi.org/10.2172/978356>.
- . 2010. "Options Study - Phase II." INL/EXT-10-20439. Idaho Falls, ID: Idaho National Laboratory. <https://doi.org/10.2172/1009149>.

- Wigeland, Roald A., Theodore H. Bauer, Thomas H. Fanning, and Edgar E. Morris. 2006. "Separations and Transmutation Criteria to Improve Utilization of a Geologic Repository." *Nuclear Technology* 154: 95–106. <https://doi.org/10.13182/NT06-3>.
- Williams, P.M., T.L. King, and J.N. Wilson. 1989. "Draft Preapplication Safety Evaluation Report for the Modular High-Temperature Gas-Cooled Reactor." NUREG-1338. Washington, DC: NRC (Nuclear Regulatory Commission). <https://www.nrc.gov/docs/ML0527/ML052780497.pdf>.
- Windes, W, G. Strydom, R. Smith, and J. Kane. 2014. "Role of Nuclear Grade Graphite in Controlling Oxidation in Modular HTGRs." INL/EXT-14-31720, Revision 0. Idaho Falls, ID: Idaho National Laboratory.
- WNA (World Nuclear Association). n.d. "How Can Nuclear Combat Climate Change?" World Nuclear Association. n.d. <https://www.world-nuclear.org/nuclear-essentials/how-can-nuclear-combat-climate-change.aspx>.
- . 2020. "Economics of Nuclear Power." WNA (World Nuclear Association). 2020. <https://www.world-nuclear.org/information-library/economic-aspects/economics-of-nuclear-power.aspx>.
- . 2021. "Nuclear Power in India." WNA (World Nuclear Association). 2021. <https://www.world-nuclear.org/information-library/country-profiles/countries-g-n/india.aspx>.
- X-Energy. n.d. "We Are a Nuclear Reactor & Fuel Design Engineering Company." X-Energy. n.d. <https://x-energy.com/nuclear-reactors>.
- Xia, Bing, Fu Li, Chunlin Wei, and Fubing Chen. 2014. "Neutronics and Thermal-Hydraulics Study on Maximizing In Situ Utilization of U-233 in the Th-U5-Fueled HTR-PM." In *Proceedings of NENE 2014*. Portoroz, Slovenia: Institute of Nuclear and New Energy Technology, Tsinghua University. https://arhiv.djs.si/proc/nene2014/pdf/NENE2014_402.pdf.
- Xu, Zhiwen. 2003. "Design Strategies for Optimizing High Burnup Fuel in Pressurized Water Reactors." PhD dissertation, Cambridge, MA: Massachusetts Institute of Technology. <https://dspace.mit.edu/handle/1721.1/16603>.
- Youinou, Gilles J., Samuel E. Bays, and Giuseppe Palmiotti. 2020. "Scoping Analysis of Sodium Cooled Fast Spectrum Test Reactor Cores." United States: Idaho National Laboratory. <https://doi.org/10.2172/1598332>.
- Yu, C., J. Wu, C. Zou, X. Cai, Y. Ma, and J. Chen. 2018. "Thorium Utilization in a Small Modular Molten Salt Reactor With Progressive Fuel Cycle Modes." *International Journal of Energy Research* 43: 3628–39. <https://doi.org/10.1002/er.4511>.
- Zanolli, Lauren. 2015. "A Radical Plan to Power the Planet for 72 Years." *Wired*, October 2015. <https://www.wired.com/brandlab/2015/10/a-radical-plan-to-power-the-planet-for-72-years>.
- Zarif, Javad. 2020. January 5, 2020. <https://twitter.com/JZarif/status/1213900666164432900>.
- Zhang, Hongbin, Haihua Zhao, Vincent Mousseau, and Ronaldo Szilard. 2009. "Design Considerations for Economically Competitive Sodium Cooled Fast Reactors." In *Proceedings of ICAPP '09*. Tokyo, Japan. https://www.researchgate.net/publication/228705398_Design_Considerations_for_Economically_Competitive_Sodium_Cooled_Fast_Reactors.
- Zou, C.Y., C.Z. Cai, C.G. Yu, J.H. Wu, and J.G. Chen. 2018. "Transition to Thorium Fuel Cycle For TMSR." *Nuclear Engineering and Design* 330: 420–28. <https://doi.org/10.1016/j.nucengdes.2018.01.033>.
- Zwartsenberg, Gijs. 2016. "SAMOFAR—Why the Slow Coming of a Fast Reactor Actually Speeds up the Development of 'Slow' MSR's." *Thorium MSR Foundation*, July 5, 2016. <https://articles.thmsr.nl/samofar-why-the-slow-coming-of-a-fast-reactor-actually-speeds-up-the-development-of-slow-msr-s-e20b32ab3341>.

“Advanced” Isn’t Always Better

Assessing the Safety, Security, and Environmental Impacts of Non-Light-Water Nuclear Reactors

Given the urgency of the climate crisis, rigorous evaluation is needed to avoid wasting time or resources in the pursuit of high-risk energy concepts.

If nuclear power is to play an expanded role in helping to mitigate climate change, newly built reactors must be demonstrably safer, more secure, and more economical than current generation reactors. One approach to improving nuclear power has been to pursue the development of non-light-water nuclear reactors, which differ fundamentally from today’s light-water-reactors. But is different actually better? The answer is “no” for most designs considered

in this assessment comparing non-light-water reactors to light-water reactors with regard to safety and security, sustainability, and the risks of nuclear proliferation and nuclear terrorism. The study from the Union of Concerned Scientists (UCS) recommends that policymakers, private investors, and regulators fully vet the risks and benefits of these technologies before committing the vast time and resources needed to commercialize them.

**Union of
Concerned Scientists**

FIND THIS DOCUMENT ONLINE:
www.ucsusa.org/resources/advanced-isnt-always-better

The Union of Concerned Scientists puts rigorous, independent science to work to solve our planet’s most pressing problems. Joining with people across the country, we combine technical analysis and effective advocacy to create innovative, practical solutions for a healthy, safe, and sustainable future.

NATIONAL HEADQUARTERS

Two Brattle Square
Cambridge, MA 02138-3780
Phone: (617) 547-5552
Fax: (617) 864-9405

WASHINGTON, DC, OFFICE

1825 K St. NW, Suite 800
Washington, DC 20006-1232
Phone: (202) 223-6133
Fax: (202) 223-6162

WEST COAST OFFICE

500 12th St., Suite 340
Oakland, CA 94607-4087
Phone: (510) 843-1872
Fax: (510) 843-3785

MIDWEST OFFICE

One N. LaSalle St., Suite 1904
Chicago, IL 60602-4064
Phone: (312) 578-1750
Fax: (312) 578-1751