

A Method for Developing and Testing Nuclear Reactor Systems with Space Applications and Improving Performance of Heat Pipe Reactor Systems

A Dissertation Presented for the
Doctor of Philosophy
Degree

The University of Tennessee, Knoxville

David Duff Dixon

December 2020

Copyright © by David Duff Dixon, 2020
All Rights Reserved.

First, I believe that this nation should commit itself to achieving the goal, before this decade is out, of landing a man on the moon and returning him safely to the Earth. No single space project in this period will be more impressive to mankind, or more important for the long-range exploration of space; and none will be so difficult or expensive to accomplish. We propose to accelerate the development of the appropriate lunar space craft. We propose to develop alternate liquid and solid fuel boosters, much larger than any now being developed, until certain which is superior. We propose additional funds for other engine development and for unmanned explorations—explorations which are particularly important for one purpose which this nation will never overlook: the survival of the man who first makes this daring flight. But in a very real sense, it will not be one man going to the moon—if we make this judgment affirmatively, it will be an entire nation. For all of us must work to put him there. Secondly, an additional 23 million dollars, together with 7 million dollars already available, will accelerate development of the Rover nuclear rocket. This gives promise of some day providing a means for even more exciting and ambitious exploration of space, perhaps beyond the moon, perhaps to the very end of the solar system itself... ...If we are to go only half way, or reduce our sights in the face of difficulty, in my judgment it would be better not to go at all. —John F. Kennedy May 25, 1961

Abstract

The concept of using nuclear reactor technology in propulsion systems is nearly as old as nuclear reactors themselves. Numerous publications, many at very primitive conceptual levels, suggested nuclear propulsion in the 1940s. Since then, numerous attempts have been made to build and fly a nuclear rocket, yet there have been no such systems advanced to the point of a system test in more than half a century.

This dissertation presents an approach to reactor system development, a testing approach developed by the author that resulted in the operation of the world's first heat pipe reactor system, and the world's first reactor system to produce electricity with Stirling engines. In the process of concept development, the author identified an opportunity to advance the state-of-the-art in heat pipe engineering, and presents a multi-stage heat pipe configuration that significantly improves heat pipe performance, particularly as heat pipes lengths increase.

The operation of the Demonstration Using Flattop Fission (DUFF) and the multi-stage heat pipe configurations herein significantly increase the likelihood of a heat pipe reactor system finding a place in regular use. Given the inherent safety of nuclear systems generally, and of heat pipe reactor systems over their conventional pumped-loop counterparts, there is hope that both space exploration and terrestrial energy security benefit from these designs.

Table of Contents

1	Introduction	1
2	Background	3
2.1	Historical Reactor Development	3
2.1.1	Development Approaches	7
2.1.2	Past Space Nuclear Power System Development Efforts	8
2.1.3	Historical Nuclear Propulsion Efforts	14
2.2	Principles of Heat Pipe Operation	17
2.2.1	Heat Pipe Reliability	22
2.3	Principles of Rocket Operation and Evaluation	26
2.3.1	Nuclear Rocket Engine System Operation	27
2.3.2	Nuclear Electric Propulsion	29
3	Reactor Development Plan Targeting Low-Enrichment Nuclear Rockets	30
3.1	Development Plan Overview	31
3.2	Sub-Kilowatt Nuclear Demonstration	32
3.3	Sub-Kilowatt Non-Nuclear and Nuclear Tests	32
3.4	Phase II: Fission Surface Power (FSP)-Derived Test	34
3.4.1	Facility Needs, Licensing Process, Estimated Facility Costs	36
3.4.2	Conceptual Design	36
3.4.3	Design Process	37
3.5	Megawatt-Class Terrestrial Reactor Test	37
3.6	Evaluation of Nuclear Thermal or Nuclear Electric for Downselect	38

3.7	Large-Scale Nuclear Propulsion Test	39
4	Megawatt-Class Terrestrial Reactor Demonstration	40
4.1	Facility Needs, Licensing Process, Estimated Facility Costs	43
4.2	Scaled Non-Nuclear Testing	44
4.2.1	Required Facilities for Non-Nuclear Testing	44
5	Double-Barrier Fuel and Annular Heat Pipe Assembly	46
5.1	Introduction	46
5.1.1	Existing Heat Pipe Reactor Concepts	47
5.2	Impact on Core Layout/Neutronic Design	48
5.2.1	Redundant Barrier Between Fission Products and Atmosphere	50
5.2.2	Reduction of Likelihood of Cascading Failures	51
6	Figure-Eight Flow Path with Constant Wick Diameter	55
6.1	Design Constraints	55
6.2	Physical Assembly and Preventing Entrainment in Crossover Region	57
6.2.1	Drawbacks	58
6.3	Straight Concentric Annular Heat Pipe with Diametrical Wicks	59
7	Multi-Stage Heat Pipe Combination	61
7.1	Inverted Multi-Stage Configuration	62
7.2	Design Process	64
7.2.1	Comparing Multi-Stage Heat Pipe Performance to Conventional Heat Pipe	67
7.3	Design Features and Trades	75
7.3.1	Effective Operating Temperature	78
7.3.2	Freedom to Adjust Heat Transfer Capacity for a Given Fuel Element	78
7.3.3	Power Peaking	79
8	The Demonstration Using Flattop Fission (DUFF) Experiment	80
8.1	Purpose and Role of Duff Experiment	81

8.2	Developing the Reactor Concept Associated with Nuclear Test and its Application in Space	82
8.2.1	Design and Analysis of the Simple Fission Reactor	85
8.3	Identifying the Opportunity for a Nuclear Test	88
8.4	Assembling a Team: Strategic Partners, Organizations, and Individuals . . .	89
8.4.1	Necessary Team Qualifications	89
8.4.2	Necessary Organizations	90
8.4.3	Strategic Exclusion of Certain Entities	90
8.4.4	Getting Started, Finding Funding	92
8.5	Advancing the Design, and Preparing for the Test	94
8.5.1	Identifying Candidate Experiments	95
8.5.2	Heat Pipe Selection/Design	95
8.5.3	Facilities and Experiment Preparation	98
8.5.4	Test Assembly and Operation	100
9	Conclusions	105
	Bibliography	108
	Appendix	114
A	Nuclear Thermal Rocket Methods Development	115
A	Prior Art	115
B	Integrated Thermal and Neutronic Calculations in NTRgen	121
B.1	Channel Flow Rate Calculation	123
B.2	Shutdown Model	125
C	Shape-Factors Correlating Finite-Difference Analysis to 1-D Heat Transfer Calculations	127
Vita		130

List of Tables

2.1	Summary of Space Nuclear System Development Efforts as presented herein, and Estimated Completion Costs, not adjusted for inflation.	17
7.1	Minimum dimensions used for unit cell comparison of annular in-core heat pipe.	77

List of Figures

2.1	Generalized Heat Pipe Flow Schematic	18
2.2	Representative nuclear thermal rocket flow schematic, Small Engine concept.	29
4.1	Segmented core concept (top), with individual core modules connected to each other via a common heat exchanger (containing many more brazed/welded joints) as incorporated in (30). Proposed concept (above) has only electrical interconnections between modules, allowing for more rapid deployment and transport in smaller packages if necessary. This approach also reduces the number of heat pipe integration operations per core block.	42
5.1	Conventional Heat Pipe and Fuel Layout with Discrete Fuel and Structural Block/Monolith and 2:1 Fuel Pin to Heat Pipe Ratio	49
5.2	Combined Heat Pipe and Fuel Unit Cell Equivalent	49
5.3	Conventional Heat Pipe Reactor Failure showing Power Delivery to Adjacent Heat Pipes	52
5.4	Annular Heat Pipe and Fuel Unit Failure showing Power Delivery to Adjacent Heat Pipes	53
6.1	Flow diagram showing figure eight heat pipe. Not to scale; ratio of heat pipe length to diameter much greater than shown.	56
6.2	Diagram showing crossover region of figure eight heat pipe.	56
6.3	Flow diagram showing straight annular heat pipe. Diametrical wicks shown in cross-sectional view. Not to scale; ratio of heat pipe length to diameter much greater than shown.	59

6.4	Flow diagram showing straight annular heat pipe. Diametrical wicks shown in cross-sectional view. Not to scale; ratio of heat pipe length to diameter much greater than shown.	60
7.1	Flow diagram showing straight annular heat pipe. Diametrical wicks shown in cross-sectional view. Not to scale; ratio of heat pipe length to diameter much greater than shown.	62
7.2	Flow diagram showing multi-stage heat pipe in conventional heat pipe concept core block. Not to scale; ratio of heat pipe length to diameter much greater than shown.	63
7.3	Diagram showing dimensional parameters for multi-stage fuel-in-heat-pipe element. Not to scale; ratio of heat pipe length to diameter much greater than shown.	64
7.4	Schematic for comparison of conventional and multi-stage heat pipe.	68
7.5	Ratio of multi-stage heat pipe limiting power to equivalent single-stage heat pipe for a specific point design as a function of length, maintaining in-core and condenser parameters and cross-sectional geometries.	70
7.6	Limiting power for multi-stage and equivalent conventional heat pipes for a specific point design. Note that while the power ratio becomes constant, total power carried by either heat pipe type declines as length increases.	71
7.7	Fraction of unconstrained heat pipe length assigned to annular heat pipe segment. Plot shown for sodium heat pipe with sintered wick, 1 cm diameter at 900K, with fixed 0.5m condenser (heat rejection to power conversion) and evaporator (in-core).	72
7.8	Limiting power as a function of temperature for a conventional heat pipe (red), and its multi-stage equivalent (blue). The annular (dashed grey) and in-core (yellow) heat pipe segments are also shown, but total power is limited by the condenser section in this case.	73

7.9	Limiting power as a function of temperature for a conventional heat pipe (red), and its multi-stage equivalent (blue), for a similar heat pipe with a thinner wick (i.e. less capillary pumping, greater flow area). In this case, the in-core segment is limiting for the multi-stage heat pipe.	74
7.10	Unit cells for conventional (left) and annular (right) heat pipe configurations.	75
8.1	^{238}Pu allocations in 2013.(19)	81
8.2	Initial Simple Fission Reactor schematic/design developed prior to DUFF test, shown alongside Advanced Stirling Radioisotope Generator concept for scale.	83
8.3	Initial Simple Fission Reactor schematic/design developed prior to DUFF test, showing components of design.	84
8.4	Flattop Schematic Drawing.	96
8.5	Heat pipe testing at Glenn Research Center, 2008.	97
8.6	First operation of Buzz converters on test stand, 90 days prior to experiment	97
8.7	Buzz Converters with Copper Heat Pipe Mounting Block	98
8.8	Flattop with Heat Pipe Inserted (end of heat pipe visible as small silver circle in center of assembly)	101
8.9	DUFF team in control room on September 13, after first DUFF run.	102
8.10	DUFF setup and attendees for second DUFF run.	103
8.11	Reactor power, temperatures, and Stirling engine output from DUFF first run.	103
8.12	Reactor power, temperatures, and Stirling engine output from DUFF second run.	104
A.1	Prior Nuclear Rocket Design Process Overview	117
A.2	TMSS Fuel Element Analysis	118
A.3	TMSS Tie Tube Analysis	119
A.4	TMSS Zoning from 2006 LANL-MSFC Effort	120
A.5	Flow Rate Solver	121
A.6	Coolant Hole Diameter Calculation	122
A.7	Fuel temperatures with only radiative cooling as a function of decay power. Note that decay power falls to 1% very quickly after shutdown.	127

A.8 Thermal Desktop Shape Factor Calculations	128
---	-----

Chapter 1

Introduction

The concept of using nuclear reactor technology in propulsion systems is nearly as old as nuclear reactors themselves. Numerous publications, many at very primitive conceptual levels, suggested nuclear propulsion in the 1940s. Since then, numerous attempts have been made to build and fly a nuclear rocket, yet there have been no such systems advanced to the point of a system test in more than half a century.

It's not generally for a lack of interest or political will that these systems have not been successful. In his famous 1961 address, President Kennedy specifically called out nuclear rocket development, right alongside the announcement of the moon challenge. Development of a nuclear rocket, particularly a nuclear thermal rocket, is an inherently challenging task. Large temperature gradients, high temperatures, difficult operating conditions, and the coupled nuclear and thermal effects combine to make design and testing, particularly non-nuclear testing, very difficult. Add to that the political challenges of an expensive development program, all the risks of launching a nuclear system, and the relative lack of cross-disciplinary experience among nuclear and rocket engineers and it should be no surprise that such a system has not flown.

This dissertation set out to challenge the idea that it cannot be done, and present a logical approach, conceptual designs, and the methods used to generate those designs, that could lead to a real system should the political will and funding exist. While not initially the intent, this dissertation also introduces fundamentally new heat pipe reactor configurations, promising increased fault tolerance and greater design flexibility. The concept and plan are

designed to take risk, first and foremost, by testing real systems, but to reduce that risk by taking it in measured, hopefully manageable pieces. As a result, and reflecting the realities of building new nuclear systems, the focus is as much on the political strategy and regulatory framework as it is on the technical aspects of the system. The operation of the world's first heat pipe reactor demonstrates that engineering decisions can solve political problems, and vice-versa. Many of the solutions to longstanding problems impeding reactor development are not complex, but require using a set of tools—from design feature to policy loophole—in a way that was not obvious prior to doing so.

Ultimately, the goal here is to demonstrate that it is possible to both develop and operate nuclear systems intended for space applications on the ground, and to put forward an approach that breaks from tradition of recent decades. It is also the goal to establish a path that, even if not executed to its end, still advances the state-of-the-art in compact nuclear reactor systems beyond the realm of studies on paper.

Chapter 2

Background

While a long list of reactor concepts and designs have been presented over the last three-quarters of a century, only a small fraction of those make it to reality. This is both a result of the nature of research and development generally, and the complex nature of both the technologies involved and the political (including regulatory) obstacles that exist.

Most in the nuclear industry are familiar with the historical milestones associated with the “firsts”, including Chicago Pile-1, X-10, EBR-I, and others. Less familiar is the discussion of how, and why, those systems were successful, and others were not. Understanding what has made reactor development efforts successful or not is critical to properly structuring a new reactor development effort so as to maximize (but of course, not guarantee) success in getting not just to a conceptual design or test reactor, but eventually to a deployable system, whether for electricity production, propulsion, or other useful purpose.

2.1 Historical Reactor Development

For the purposes of discussion, it is useful to segregate the history of reactor development for production of nuclear materials and defense purposes from that of reactor development for the purposes of energy production, whether that energy is used for electricity production, process heat, or in the case of nuclear propulsion systems, thrust. In the years following World War II, and accelerated with Atoms for Peace, the number of reactor designs that existed at a conceptual level, as well as the number in various stages of design and construction, increased

rapidly. When you look at so-called “advanced” reactor concepts today, and examine them in terms of their fundamental technologies and fuel systems, most trace their roots to concepts introduced in the late-1940s and 1950s.⁽⁵⁾ This includes light and heavy water reactors largely as used today, but also gas-cooled, liquid metal, molten salt, and aqueous reactor technology. Each of these technologies has been used in at least one or more actual reactor systems, of varying degrees of technical maturity, and all of them implemented at least once in a critical system by the 1960s.

For commercial applications, the world has relied heavily on water-moderated, water-cooled reactor systems, with graphite-moderated, water-cooled making up a significant, but much smaller contribution to the nuclear energy technology portfolio. While these basic technologies are certainly not without issue, they benefited from industrial familiarity (boiler and steam turbomachinery), relatively inexpensive and readily available materials, dynamic stability (i.e. negative temperature coefficients, good cooling properties, easy pumping, and a lack of freezing at normal temperatures). The benefit of hindsight makes water-cooled, water-moderated reactors, particularly light-water reactors, seem like an obvious technology to adopt.

On the other hand, what would today be considered advanced designs, such as liquid metal-cooled systems, were in many cases as mature, or more mature, than their water-cooled counterparts. By the early 1950s, both mercury and sodium-potassium reactor systems had been operated at temperatures high enough to prove their viability for large-scale power production. Gas-cooled reactor systems were less mature, mostly lagging not in system concept maturity but more so the availability of suitable materials, particularly for the fuel system. While it’s ultimately a gas-cooled reactor, many of the challenges for gas-cooled reactor systems in stationary power applications were less relevant for nuclear propulsion systems.

Standing in stark contrast to each of these systems, each with a relatively mature set of concepts, practical designs, and test systems/operating history, the idea of a heatpipe-cooled reactor is much more recent. The first patent granted for what would today be called a heat pipe, 2,350,348,⁽¹⁶⁾ was not granted until 1942, and the more popular attribution of the heat pipe’s invention, to George Grover, was not until 1963—well after most of the aforementioned

reactor system types had been tested. (Grover is also credited by many with the first notion of using heat pipes in nuclear reactor applications.)

It is the relative novelty of the heatpipe-cooled reactor system, and its successful development and nuclear testing in the 21st century, that will allow a comparison of successful modern development with successful historical development, along with comparison of the successful modern development with many shelved attempts to develop and advance new (or more broadly, simply non-water) reactor types.

Regulatory Framework

As we examine a number of prior nuclear reactor development efforts, it is necessary to consider the regulatory framework in which they occurred. Obviously, regulation can be changed, but that change is extremely expensive and risky in the context of program development. We must acknowledge that while some political action is justified and necessary (e.g. changes to regulations not applicable to a certain system type, funding authorization), other governing law, regulatory bodies, and political constraints should (and likely will) continue in place.

As it was largely temporary, and not likely to reoccur, we defer for now the earliest contribution to regulation that occurred during World War II. Moving ahead to the immediate aftermath of the wartime effort, nuclear material was still government-owned, and government-controlled in the United States. Congress started debating how to regulate this new atomic energy, wanting to display appropriate use in both war-related and peacetime endeavors. Almost a full year after the bombings in Japan, on August 1, 1946,⁽⁷⁾ President Truman signed the Atomic Energy act, which established the Atomic Energy Commission (AEC) and moved control of the atom out of the War Department.

This new agency was directed by civilians, but under strict government control. The government owned all of the fissionable materials (Pub. L. 79-585, Section 5(a)(2)) and the production plants for fissionable materials (Pub. L. 79-585, Section 4(b)). The facilities where those materials could be used required license (Pub. L. 79-585, Section 7(a)) from the government, and any material used remained government property. The agency's main

purpose remained military. It was, however, encouraged to look into peaceful uses of nuclear technology.

Nevertheless, peaceful purposes and opportunities for power generation were advanced under this framework. The AEC and its contractors had, by 1951, assembled technical data, design information, and conceptual drawings for each of the reactor types previously outlined. A number of these had already advanced to operational status by that time, including Clementine(6) and EBR-I.(28)

Of course, most of those reactors were built specifically for scientific research, not energy production needs. Many were focused on gathering physics data, including nuclear cross sections and characterizing heat transfer. Others were built for fuel cycle research, principally regarding production of plutonium for military use. They were designed and built for a relatively short mission—to operate long enough to complete the intended research.

Safety analysis under the AEC was also quite different than under today’s Department of Energy or Nuclear Regulatory Commission, or any of the other regulatory bodies that might be involved in a nuclear project. Engineering judgement was used to determine what was safe enough or needed further study before proceeding—for many tasks, there was no empirical standard by which risks could be quantified. Workplace safety generally, environmental protections, and hazards of radiation exposure were viewed and regulated quite differently than today. OSHA, EPA, NRC and DNFSB were all formed decades later.

As it stands today, any domestic non-military reactor operations must fit into one of two regulatory frameworks: that of the Department of Energy, operating its own nuclear facilities, or of the Nuclear Regulatory Commission (NRC), which regulates both commercial nuclear power facilities and research reactors of the sort found on many university campuses. The NRC’s regulatory authority also extends to new Department of Energy demonstration reactors when operated “for the purpose of demonstrating the suitability for commercial application of such a reactor.” The Department of Energy’s regulatory authority under federal law is broadly contained at 42 USC 2201(i)(3) and 42 USC 2051 with regard to safety and operation of its own facilities, while a variety of functions of the former Atomic Energy Commission were transferred to the Nuclear Regulatory Commission 42 U.S. Code §5841, part of the Energy Reorganization Act of 1974. In the proposed plan for nuclear

reactor development, it is important to note that the intended applications are somewhat narrowly construed, and limited to government owned and operated space power systems (i.e. NASA), and those intended for use by the Department of Defense.

Since the passage of the Energy Reorganization Act of 1974, no new nuclear reactor facilities have been constructed and operated other than those already contracted prior to passage. The NRC has certainly granted licenses for construction and operation, but each of the licensees that has operated commenced the project prior to the reorganization. Today, there are a small number of concepts in early stages prior to licensing, one combined construction and operating license application submitted prior to construction, and two large commercial reactors currently under construction, and there is hope that these will be completed even if significantly over-budget and behind schedule. The Department of Energy, on the other hand, operates a fleet of nuclear reactors (including critical experiment machines) for which it assumed ownership and control during the reorganization. It had never, prior to DUFF,¹ operated a nuclear system that produced electricity.² Even though both DUFF and KRUSTY(39) produced electricity, they nevertheless relied on critical experiment machines with operational history and pre-existing licensure.

In terms of launching space fission power systems, a fairly significant change occurred in late-2019, with issuance of a new presidential memorandum on the launch approval process. Under the prior framework, presidential approval was required for the launch of *any* nuclear system with the potential for criticality. The updated framework still requires presidential approval for some launches, including any government fission system containing fuel other than low-enriched uranium. Under the current memorandum, the launch authorization authority for LEU fission systems is the head of the sponsoring agency.(47)

2.1.1 Development Approaches

It hardly needs mentioned here that the environment in which a reactor advances from early concept to operational system is in many ways far different today than it was at the dawn

¹Demonstration Using Flatop Fission. See Chapter 8

²Excluding Naval Reactors, which reports to both the Department of Energy and the Department of the Navy

of the atomic age. In many respects, at the dawn of the atomic age and well into the 1950s, nuclear development occurred as if in the old wild west: limited knowledge, largely unregulated (relative to today), more hazardous, less environmental concern, and virtually non-existent computing capability. The demands on the industry to advance technology were also much greater, especially as political pressure demanded practical peacetime uses keep pace with the arms races that threatened to end the world.

If we were to discuss here only those domestic systems tested in some form to-date, they essentially fall into 3 groups: ROVER/NERVA, SNAP, and DUFF/KRUSTY. Limiting the discussion further, to domestic systems actually flown in space, the list shortens to only one system, SNAP-10A. This is despite significant efforts in the meantime, including large space reactor development efforts including SP-100 and Project Prometheus, along with a never-ending series of smaller programs, including the Affordable Fission Surface Power System, the U.S. TOPAZ program, several propulsion efforts, and countless studies for various mission concepts from ice melters on Jupiter to other Radioisotope Power System replacement concepts.

Earlier programs, including ROVER/NERVA and SNAP, fell under the auspices of the Atomic Energy Commission. Both were supporting military objectives, with large budgets, and clear purpose. For much of the last half-century, however, reactor development efforts have followed an all-too-common cycle, managed by the Department of Energy: assemble a committee to evaluate a set of pre-conceptual designs for a particular mission or class of missions that do not currently have a viable power source, investigate the problem for several years, and when the program either grows too large or the flagship mission changes, shelve the program.

2.1.2 Past Space Nuclear Power System Development Efforts

Systems for Nuclear Auxiliary Power (SNAP)

As a broader program, Systems for Nuclear Auxiliary Power, or SNAP, was largely successful in accomplishing its stated goals. SNAP developed a fairly significant number of radioisotope

power systems, and laid the groundwork for the current radioisotope power system technology still in use today.

Its successes in reactor development are somewhat mixed, depending on the evaluation criteria. SNAP’s ground tests, and successful launch of a reactor into space certainly stand as far greater successes than any program since. Prior to SNAPSHOT/SNAP-10A’s launch in 1965, the SNAP program operated a series of reactors, and developed a safety framework for reactor launches. While first-of-a-kind, it was a product of a different era, and a reactor like SNAP-10A will likely never be launched again.

Unlike generally accepted criteria for launch of a nuclear system today, SNAP was expected to go critical in a number of launch accidents. To assess what could happen in a launch accident (i.e. prompt criticality, energy release, and dispersion of fissile, radioactive, and hazardous material), a series of transient tests were performed recreating those postulated accidents.

The worst-case accident considered was an aborted launch resulting in the reactor crashing into the ocean, resulting in the reactor going prompt critical, and exploding. Instead of designing to prevent such an event, an experiment was conducted that actually *caused* the worst-case version of that to occur, sitting on a rail car in Idaho, on April 1, 1964. The objective was to measure the radiological release, and show that it would not be a significant hazard.

To be clear, this was a test reactor immersed in a tank of water, outdoors, configured such that when a charge of black powder was ignited, a cylinder containing B_4C was ejected causing a reactivity insertion of nearly $\$4^3$ inserted. The energy released was equivalent to approximately 10 kg of TNT, with peak power estimated at 18 GW.(23) Such a test would be unconscionable today, and would likely violate a number of international treaties.

The program was successful in getting one reactor into space safely, but despite all development efforts, the system failed after 43 days. The failure is commonly reported as caused by a faulty “voltage regulator”, though the actual failure is somewhat more complex:

³The excess reactivity limit for the DUFF experiment of reactivity was 0.80. Inserting only $\$1$ of excess reactivity means a critical system becomes *prompt critical*.

The telemetry data revealed a vehicle status inconsistent with any logical mode of operation of the malfunction and failure detection system, failure of a number of independent electrical components and a number of operations which could only be explained by ground commands that had not been sent.(51)

Remnants of the SNAPSHOT system, containing the SNAP-10A reactor system remain in orbit. It continues to shed debris at regular intervals, which started prior to 1979.(46)

Some remnants of SNAP system concepts, especially the proposed use of uranium nitride fuel and lithium coolant, would resurface in the SP-100 program. Advances in safety strategy (and public expectations) likely mean that any future system with similar objectives would share little technology with SNAP-10A.

SP-100

The SP-100 program was launched in 1983 as a “triagency” project involving DARPA, DOE’s Office of Nuclear Energy, and NASA’s Office of Aeronautics and Space Technology, coupled with the Strategic Defense Initiative. At its inception, the following statement presents a posture that would be evident throughout its 12-year duration:

A space nuclear reactor power system based on already proven technology is ready for a ground test phase now, but its performance would be limited in terms of power-to-mass ratio, total power, and life. It is likely that potential users would not find it interesting enough to employ in their missions.(48)

SP-100 advanced a uranium nitride reactor system using pumped lithium coolant and thermoelectric power conversion. At $100kW_e$, maximum mass of 3,000 kg, and occupying no more than one-third of the space shuttle’s cargo bay, the design requirements drove the need for the very high temperature system.

The program initially planned a ground nuclear test that would occur in 1986, which by 1995 (when the program was canceled) was projected to occur in 1998. At that time, the program had spent more than \$500 million. No prototypes were tested, no systems flown.

Congressional testimony in 1992(17) points out that projected cost estimates had more than tripled, and at less than 9 years in, the program was 13 years behind schedule, and

costs to complete phase II (ground-based component testing) were estimated at over \$2 billion through 2009.

TOPAZ (U.S.)

As the Soviet Union collapsed in 1989, the U.S. was offered a pair of unfueled TOPAZ reactor systems. Over several years, plans to fuel and fly the systems were abandoned, along with efforts for flight qualification of a derivative thermionic reactor system.

By itself, U.S. TOPAZ program was fairly unremarkable: some non-nuclear testing of a system that had already been designed and built. But it played a role in the decision to cancel the SP-100 program, with the Department of Defense arguing that development and space qualification of a thermionic system based on the purchased TOPAZ systems would result in a four-fold reduction in cost (from $> \$2B$ to $< \$500M$) and deliver a system ready for launch in seven years instead of 17.⁽¹⁷⁾ It certainly sounded credible: by 1989, six ground nuclear system tests of TOPAZ reactors had already been completed.

The Department of Energy argued the estimates were not realistic, and that splitting already limited funds would reduce the likelihood of success. Nevertheless, two systems were purchased and delivered in May 1992.⁽⁴⁴⁾ Non-nuclear testing was performed over a period of approximately three years (ending in early 1995); cancellation of the program followed shortly thereafter.

Jupiter Icy Moons Orbiter/Project Prometheus

Following the SP-100 and TOPAZ programs, there was a fairly significant gap in funding for space reactor development, similar to the gap between SNAP and SP-100. Like SP-100, Project Prometheus was announced 2 years into a presidential administration, and grew quickly. Unlike SP-100, Prometheus seemingly had an “all things nuclear” umbrella, with “key program components” including everything from radioisotope generators to multimegawatt power and propulsion systems.⁽³⁴⁾

Both the Multi-Mission Radioisotope Thermoelectric Generator (MMRTG) and Stirling Radioisotope Generator (SRG) were a part of Prometheus, with an MMRTG having flown on Mars rover, *Curiosity*. The Advanced SRG program continued beyond Project Prometheus,

but was terminated in 2013. While unsuccessful in the effort to qualify a Stirling-based radioisotope system, there is no question that despite limited ^{238}Pu supplies that the RTG capability is more robust today than in 2003 at the program’s launch.

The flagship piece of Project Prometheus was an electric propulsion system driven by a high-performance 200kW_e nuclear reactor. Over the course of less than 3 years, the program spent over \$500M performing concept studies, trying to reestablish fuel fabrication capabilities from SP-100, marching toward a ground test with a \$1B cost estimate. The concept development was significant, and was unique among the programs discussed so far in that DOE-Naval Reactors was heavily involved in the design studies.

Like those before it, cost estimates grew significantly as the project progressed. When the Jupiter Icy Moons Orbiter mission was shelved, the need for a reactor to power it also vanished. With the shift in priorities to human exploration in Project Constellation, there was an effort to shift space nuclear development to stationary power systems for the surface of the moon or Mars.

Part of that effort at the end of Project Prometheus involved concept studies resurrecting UZrH fuel, Hastelloy-X cladding, and pumped NaK from the SNAP era.(10) Despite the push to use materials and approaches with “heritage,” their relative obscurity and dormancy for nearly half a century presented its own obstacles.

Affordable Fission Surface Power System

Over the next five years, Fission Surface Power dominated the space nuclear system landscape, albeit at very low funding levels compared to Prometheus. The development team promoted the use of less-aggressive performance parameters, non-nuclear component testing, and nuclear modeling such that the system could be launched without a ground nuclear test. The argument was sound at the outset: if the cost of a ground nuclear test would exceed the cost of launching a system and testing it in space, what is the objective of the ground test?

While heavily focused on affordability, the program still suffered from a multi-lab structure under DOE-NE management. A number of design decisions clearly reflect the idea that each of the national labs was vying for its share of the pie. Two deserve mention

here: the use of an annular linear induction pump (ALIP) to move the NaK coolant, and the use of two parallel intermediate loops linking the primary reactor coolant loop to a set of Stirling converters. A reliability model showed that no matter how reliable or unreliable a given component, the configuration with the intermediate heat exchanger was less reliable. It had a set of possible failures where the system could operate at half power, however, and that seemed to be more important, even though the odds of complete failure were still higher than with the direct loop configuration studied.

Idaho National Laboratory had primary design responsibility for the ALIPs, a pump technology that is simple in concept with no moving parts, but difficult to perfect with any reasonable efficiency. Oak Ridge National Laboratory was responsible for heat exchanger design, a rather straightforward effort but one not needed at all if the coolant loop passing through the reactor delivered that coolant directly to the Stirling converter heads. The direct configuration meant one ALIP and no intermediate heat exchanger, the intermediate loop configuration meant three ALIPs (two different designs) and an intermediate heat exchanger. The loop configuration downselect significantly increased the project cost, but assured the involvement of the two labs' teams.

As cost estimates became clearer, this affordable system, even without a ground nuclear test, had still managed to grow into a billion-dollar effort. With a new presidential administration and a shift from returning to the Moon to visiting an asteroid, the end of the Constellation project effectively also meant the end of the Fission Surface Power effort. Pieces of the program would continue, including a Technology Demonstration Unit built at NASA-GRC, but the reactor development and progress toward a flight system was halted in 2010.

Small Fission Power System Feasibility Study

The Small Fission Power System Feasibility Study was a very short (6 weeks) look into the idea of developing a very small reactor system for unspecified science missions. The team included a number of past space nuclear power program participants, including most of the AFSPS team. The team defined its own mission requirements ($1kW_e$, 2020

launch availability, and 15-year lifetime), and generally had a straightforward development approach, using well-known materials and technologies and reasonable operating parameters.

In the context of this dissertation, this study missed a huge opportunity. Stirling engine power conversion, sub-kilowatt power levels, and testing using the critical experiment facilities at the Nevada National Security Site were all mentioned in the study. But the study concluded that there was little merit in lower power, as “the reactor is already near the minimum size needed to sustain a fission reaction.” It selected thermoelectrics over Stirling engines for power conversion, driving the thermal power up but eliminating moving parts. It suggested component irradiation experiments at existing nuclear facilities, and a series of tests including zero-power critical test at the end of the development phase, electrically-heated system testing in the engineering phase, and zero-power “acceptance” testing during flight qualification. It did not propose a ground nuclear test as part of any of those phases of development. The final report was released in May, 2010, and sat idle.

Soviet/Russian Reactors

For completeness, it should be noted that the former Soviet Union designed, tested, launched, and crashed a series of reactors. Some, like TOPAZ, have directly influenced US development efforts. Each of the technologies used has been studied or developed at some point domestically. In terms of regulatory and political approaches, none are particularly relevant to a new system today.

2.1.3 Historical Nuclear Propulsion Efforts

Project Rover

Project Rover was the first, and the only large-scale, nuclear thermal rocket development program. Initiated in the mid-1950s, and continued essentially until 1968, the program tested 23 nuclear thermal reactors at the Nevada Test Site. The program spent approximately \$1.5B dollars (\$12B in 2020 dollars).⁽¹⁸⁾⁽²¹⁾

Performance of the test reactors was significantly improved over the life of the program but each of the first 9 engines had catastrophic failures including fuel melt, hydrogen fires,

and ejection of large fuel masses during the first 5 minutes of operation.(14) Even with progress through the PHOEBUS and NRX series of engines, by the end of the program a substantial development effort was still needed prior to actually flying a system. Some technology from Project Rover overlapped with the Aircraft Nuclear Propulsion program, but neither survived to completion.(3)

For over 40 years, Project Rover laid claim to the highest-power steady-state operation of any reactor in the world.⁴ As a research program, the advancements in high temperature materials and reactor design cannot be overlooked.

Nuclear Engine for Rocket Vehicle Applications (NERVA)

The NERVA program is often mistaken for one independent of Project Rover, but in fact it was a component of it focused on technology demonstration. Most of the subsequent (modern) work on nuclear thermal propulsion systems have relied in large part on the basic design of the NERVA engines.

Until the most recent effort, the majority retained the 19-hole fuel element geometry, basic tie-tube structure, and external reflector control. Eight tests within Project Rover were NERVA demonstration engines.

Project Timberwind and Concurrent NASA Efforts

Project Timberwind was a classified program from 1988 to 1991 under the Strategic Defense Initiative Organization (SDIO). Under SDIO management, the program spent approximately \$131 million, and was formally terminated in 1992 having spent \$196 million with another \$40 million allocated. Work on the "particle bed" concept from the program continued, with the Air Force taking over program management under the Space Nuclear Thermal Propulsion (SNTP) name. (25)

The goal of Timberwind was to "increase lift capabilities by 200 to 400 percent over current chemical rocket engines." Using the space shuttle's main engine as a reference, that

⁴At an estimated power of 4080-4500MWth, the PHOEBUS2A test was possibly the highest thermal power of any reactor until overtaken by Taishan's Unit 1 at 4590MWth, which began operation in 2018.(12)

would imply a specific impulse of 1,300 to 2,300 seconds.⁵ The program initially planned a nuclear thermal rocket launch into space in 2000.

During the same time period, NASA funded a variety of efforts looking to advance nuclear thermal propulsion. Kingsbury’s testimony in 1992 also suggests ongoing NASA expenditures on the order of \$3-4 million annually, performing concept and technology studies. Studies included solid, liquid, and gas core (40) concepts, determining that the solid core concepts derivative of the NERVA technology demonstrations were the appropriate path forward. It should be noted that at that time that plans indicated building and ground testing such a system by 2006, as part of the Space Exploration Initiative that would return man to the Moon early in the 21st century. A testing facility to support SNTP fuel elements carried an estimated cost of \$407 million to \$1.0 billion. The SNTP program was canceled in 1994, just ahead of SP-100 and TOPAZ.

Recent Nuclear Thermal Rocket Design Programs

Subsequent to the cancellation of Project Rover, there had been numerous low-level research efforts involving nuclear propulsion, but the most significant was perhaps the joint effort between NASA’s Marshall Space Flight Center and Los Alamos National Laboratory. While not much more than a one-year effort, it was sufficiently well-funded to relocate a rocket engineer to Los Alamos, where a team developed a set of nuclear and thermal analysis and design tools, and considered several design concepts. More recently, the Nuclear Cryogenic Propulsion Stage program is a much larger scale development effort, that places heavier emphasis on materials development and qualification, particularly cermet fuels, and is tasked with identifying technology challenges and programmatic issues. Whereas the 2006 effort was a small team focused on design, NCPS has made use of over 50 people, only a small portion of which were focused on engine system design. Methods arising out of both programs are discussed herein. Currently, NASA’s NTP project and DARPA’s DRACO solicitation represent a significant increase in funding for a similar NTP system, with a focus on using low-enriched uranium. While the use of LEU reduces the system’s classification to *Tier II* under the current launch approval memorandum, the size, void fraction, and moderated spectrum

⁵Specific impulse is discussed in section 2.3.

Table 2.1: Summary of Space Nuclear System Development Efforts as presented herein, and Estimated Completion Costs, not adjusted for inflation.

Program	Duration	Cost	Est. Completion Cost
SNAP	1959-1971		
ROVER	-1971	\$1.5B	
Timberwind/SNTP	1988-1994	\$250M	\$407 – \$1.0B
SP-100	1983-1995	> \$500M	\$2B
TOPAZ	1992-1995		> \$400M
JIMO	2003-2005	> \$500M	\$4B-\$8B
FSP	2006-2010	\$10Ms	> \$1B

make launch safety a formidable problem. Current approaches to ensuring subcriticality largely rely on fine wires inserted into each fuel hole, withdrawn after the system is in a stable orbit or beyond Earth orbit.

Summary

As one looks at the timelines for major space nuclear development efforts, it is impossible to ignore the boom-bust nature of funding, with major changes tending to occur roughly two years after a presidential election in which the party holding office changes. Since the Kennedy administration, these changes have occurred following all but the 1978 turnover. The reasons are not as simple as being for or against nuclear systems in space, but reflect NASA’s near-term objectives.

This suggests that for any new nuclear system plan to be successful for space applications, it must be able to be implemented within a 4 (or possibly 8) year term.

2.2 Principles of Heat Pipe Operation

In presenting the reactor design innovations in Chapter 5, it will be necessary to understand the basic principles of heat pipe operation. Heat pipes offer an extremely efficient and passive means of heat transport. The concept of a heat pipe was first described in patents 2,350,348 and 2,448,261 (Gaugler (16) and (15)). They precede the common term *heat pipe* in use today, but describe a “heat transfer device” and a “capillary heat transfer device for




Figure 2.1: Generalized Heat Pipe Flow Schematic

refrigerating apparatus” containing all of the necessary elements of a heat pipe as commonly known.

Decades later, George (20) patented a heat pipe consisting of a Niobium 1% Zirconium alloy structure, and lithium as the working fluid. Fundamentally, it described largely the same system as Gaugler, but specifically claimed a heat pipe with lithium as the working fluid. It also presented performance data using a sodium-filled heat pipe, and describes a power limit observed in testing.

In its simplest form, a heat pipe is a closed round tube containing a working fluid and a wick (which may be as simple as longitudinal grooves on the heat pipe wall) on the inside of the tube surface. If a region of the tube is heated, the working fluid will eventually begin to evaporate, increasing the partial pressure of vapor in this evaporator region. This drives a flow away from the evaporator region. Where heat is rejected from the tube, vapor condenses on the wall of the heat pipe. Capillary forces in the wick draw the condensed fluid along the wall of the heat pipe back toward the evaporator section. Gravitational forces affect both liquid and vapor fluid flows, but are generally only of significance for the liquid phase.

Generally, the functional regions of heat pipes are the (1) evaporator, where heat is added and evaporation occurs, (2) an adiabatic section, where both phases are flowing but not transferring heat out of the heat pipe, and (3) the condenser, where heat is removed from the heat pipe. These regions do not have to be in any particular order, and multiple regions of any type may be present within the same heat pipe.

While there are some gradients present, a heat pipe can be largely viewed as an isothermal chamber, with its working fluid at the saturation pressure for the operating temperature. With typical heat pipe efficiencies, the slight increase in saturation temperature (and sensible heat added as a result) and the increase in evaporator pressure are negligible.

Presume a simple heat pipe with one evaporator and condenser. The mass flow rate between them is a function of only the heat of vaporization and power being transferred.

$$\dot{m} = \frac{Q}{H_v} \quad (2.1)$$

Like other two-phase flow systems, this is why they are so much more effective than a liquid-only transport mechanism. For water operating near atmospheric pressure and temperature, heat of vaporization is on the order of 2200 kJ/kg. Therefore, a water-filled heat pipe would be able to carry 1kW of thermal energy with a mass flow rate of 0.45 g/sec, or less than 0.02 gallons (US) per minute of liquid water.

The mass flow rate in Equation 2.1 can be used along with a series of ordinary pressure drop calculations to determine the necessary capillary pumping power required to drive the circuit. The sum of all pressure drops in the flow path must equal zero, and must be less than the limiting capillary pressure.

$$\Delta P_{capillary\,limit} \geq \Delta P_{liquid} + \Delta P_{vapor} + \Delta P_{gravity} \quad (2.2)$$

For a cylindrical heat pipe, pressure drop in the vapor region is the most familiar. In the adiabatic region, the pressure drop is ordinary pressure drop in a tube, calculated with the standard fluid mechanics relationship in Equation 2.4, with a constant mass flow rate. Starting at the evaporator end of the heat pipe, vapor mass flow rate rises from zero to a maximum at the evaporator exit as heat is added, with the additional mass sourced from evaporation on the exposed wick surface. Assuming a uniform heat input over the evaporator length, the change in mass flow rate is constant, and we can approximate the pressure drop with an effective length equal to one-half of the actual evaporator length, and calculate the pressure drop as in the adiabatic region. In the condenser, pressure drop is calculated in the same manner as the evaporator assuming uniform heat removal, using an effective length equal to one half of the actual length.

$$L_{eff} = \frac{L_{cond}}{2} + L_{adiabatic} + \frac{L_{evap}}{2} \quad (2.3)$$

$$\Delta P_{vapor} = \frac{f \rho_v L_{eff} V^2}{2 D_h} \quad (2.4)$$

The pressure drop in the wick derives from Darcy's law for fluid flow in porous media:

$$\Delta P_{liquid} = \frac{\dot{m} \cdot \mu_L \cdot L_{eff}}{\rho_L \cdot K \cdot A_w} \quad (2.5)$$

K refers to the wick permeability, and A_w is the cross-sectional wick area.

Where gravitational or acceleration forces are present (i.e. when a heat pipe is operated within a gravitational field and not perpendicular to the normal force), one must also consider their contributions to the flow loop.

$$\Delta P_{gravity} = (\rho_L - \rho_V) \cdot g \sin \theta \cdot h \quad (2.6)$$

where $\theta = 0$ is horizontal, and positive values refer to an orientation where the condenser is lower than the evaporator. Note that when the evaporator is lower than the condenser, as in the case of a vertically-oriented heat pipe with the heat source at the bottom, that this gravity term decreases the wick pumping power required. Since ρ_L is much larger than ρ_V , the gravitational term is often approximated considering only the contribution from the liquid phase.

The remaining term, the capillary pumping term, is limited by a function of surface tension of the working fluid and the radii of the liquid-vapor interface in the wick structure. For screen and sintered wicks (where pores have the same dimension both parallel and perpendicular to the liquid flow), this takes the form in Equation 2.7.⁶

$$\Delta P_{capillary\ limit} = \frac{2\sigma}{r_i} \quad (2.7)$$

It is worth noting here two competing phenomena in wick design. The frictional pressure drop decreases with a higher wick permeability, while a larger pore size decreases the capillary limit. Generally speaking, wicks with larger pore sizes have higher permeability, and vice versa, with the relationship depending on the wick type (e.g. grooved, sintered, etc.).

⁶Note that r_i as used here refers to pore radius, not a heat pipe dimension.

The capillary limit can also be expressed as a maximum power, with the following form(4):

$$\dot{Q}_c = \frac{\rho_l \cdot H_{fg}}{\mu_l} \cdot \frac{K \cdot A_w}{L_{eff}} \cdot \left(\frac{2\sigma_l}{r_{eff}} - \rho_l \cdot g \cdot l_t \cdot \cos \Phi \right) \quad (2.8)$$

For a given set of operating parameters and fluid properties, this limiting power can be shown proportional to wick area, A_w , and inversely proportional to effective length, l_{eff} :

$$\dot{Q}_c \propto \frac{A_w}{L_{eff}} \quad (2.9)$$

So long as the sum of the vapor, normal/gravity, and viscous terms is less than the limiting capillary pumping term, a circulatory flow is established in the presence of a heat source and sink. If operated under conditions where (1) the vapor flow velocity is subsonic, (2) the shear forces from the vapor flowing past the wick remain low enough to avoid entrainment of the liquid phase, (3) heat fluxes are low enough that boiling within the wick does not occur, and (4) capillary flow is sufficient to maintain a fully-wetted wick in the evaporator region(s), the heat pipe will transport heat as intended, as a nearly isothermal device. This operating regime is referenced as the *heat pipe regime*.(8)

For almost all space reactor concepts to date, the heat pipes proposed have been cylindrical in shape, with some having occasional bends. In other applications, including smartphones and laptops, heat pipes often have a nearly rectangular shape, with heat input from only one side.⁷ A variety of other heat pipe configurations, including annular heat pipes, heat pipes with multiple condenser and/or evaporator regions, and oscillating heatpipes utilizing a serpentine pattern, have been demonstrated. Most commonly, wicking structures take the form of either grooved channels on the heat pipe wall, or a sintered or screen wick adjacent to (but not in contact with) the wall. The former is generally cheaper and simpler to manufacture, while the latter offers better performance. The performance gain follows not just from the wick itself but also the small gap between the wick and wall. This allows condensation at nucleation sites on the wall to develop into droplets of an increasing size, until the droplet wets and is drawn into the wick, and increases the size of the liquid-vapor interface. For liquid metal heat pipes, capillary pumping is most often the limiting

⁷Note that this geometry approaches the final geometry proposed in 6.3

phenomenon, and advanced features such as *arteries*⁸ with vents have been proposed and/or demonstrated as a means to increasing pumping limits.

Recalling that a heat pipe operates at the saturation pressure associated with its operating temperature, and that the capillary pressure is directly proportional to surface tension, certain working fluids become obvious choices for certain temperature ranges. For near-term heat pipe reactor systems operating at less than 1300K, where use of refractory metals is not necessary, the choice is almost always sodium or potassium metal. At higher temperatures ($> 1200K$), the choice shifts to lithium. At room temperature, or in space, both of these metals are in solid form. Thawing and reaching saturation is generally a fairly quick process due to the small inventory of working fluid in a heat pipe, but the phase change and sensible heat addition can be rate-limiting on startup. In particular, when vapor pressures are very low, the sonic limit is easily exceeded, and if this rate is exceeded, evaporator temperatures can exceed the nominal working temperature, taking on heat from the reactor core while only rejecting it through the small cross-sectional area of the heat pipe structure.

2.2.1 Heat Pipe Reliability

Properly designed and manufactured heat pipes are generally quite reliable. As a completely passive device, with no moving parts, they are generally at least as reliable as a pumped loop with comparable specifications. They are not immune to failure, however, with a recent NASA weather satellite (GOES-17) experiencing a loop heat pipe failure attributed to particulates in the working fluid.⁽³²⁾ While not directly suggested in the report, the most probable location for particulates to create a blockage is in the wick, where small pores could be blocked by those particulates.

While the aforementioned failure is irrecoverable, the terminology commonly used with regard to heat pipes is ambiguous at best. At the most basic level, *failure* generally references any condition where the heat pipe fails to operate as intended. As commonly used, this also includes conditions where the operating limits (boiling, capillary, entrainment, sonic) of the

⁸Arteries take the form of a small channel enclosed by wicking material, intended to permit the draining of liquid back to the evaporator region.

heat pipe are exceeded but where the heat pipe is not damaged, and where a change in operating conditions can restore proper operation.

The distinction between recoverable and irrecoverable failure conditions is particularly important for space systems and safety. The author chooses to reserve the term *failure* for irreversible conditions, and adopt the term *inoperable* to describe a heat pipe not operating in the *heat pipe regime* as intended, but which can be returned to an operable state by changing the operating conditions. A third set of limitations, where heat input does not need to be reduced to restore heat pipe regime operations, is referred to as “nonfailure limitations”.

Inoperable Heat Pipes (Recoverable Failure)

The limits of heat pipe operation, from a fluid mechanics view, have been well studied. Generally, under normal operating conditions, the capillary and boiling limits are most relevant, while sonic, viscous, and entrainment phenomena tend to limit performance while a heat pipe is heated from a cold state. A brief description of each of these limits follows.

The *viscous limit*, or *vapor pressure limit*, refers to a condition when the saturation pressure in the evaporator region is insufficient to drive flow to the condenser (i.e. the evaporator vapor pressure is equal in magnitude to the viscous forces resisting fluid flow). It is a form of *nonfailure limitation*. Generally, as temperatures increase, the viscous limit gives way to the *sonic limit* as the limiting phenomenon.

The *sonic limit* is a second nonfailure limitation that is generally encountered at on heat pipe startup. At low vapor pressures, such as would occur as a heat pipe is starting from a cold condition, vapor flow from the evaporator to the condenser must occur at very high velocities for a given mass flow rate. If the velocities are high enough, choked flow can occur. The increased pressure gradient in the vapor causes the flow loop to stall, exceeding the capillary pumping limit under those conditions. Fortunately, for both of these nonfailure limitations, the evaporator temperature begins to rise, producing more vapor. This both increases the driving pressure and reduces the vapor velocity necessary for steady-state operation. While it can be a concern in certain transient conditions, it is essentially self-correcting for a properly sized heat pipe.

The remainder of inoperable heat pipe conditions all have the same end result: dry-out of wicking material in the evaporator. When dry-out occurs, the heat pipe no longer operates in the heat pipe regime, and functions more like a cylinder filled with a static vapor. The effective conductivity of the heat pipe falls dramatically.

The *entrainment limit* refers to a condition where shear forces from the vapor phase flowing against the liquid in the wick cause droplets to break away from the wick and flow with the vapor toward the condenser. Recalling that vapor velocities fall with increases in pressure, this tends to be limiting in intermediate temperature ranges for many common fluids. Unlike the viscous and sonic limits, since this is starving the evaporator of liquid, it requires a reduction in vapor velocity (increase in temperature or decrease in power) to recover.

A limiting heat flux, frequently called the *boiling limit*, occurs where the temperature rise from the wick-vapor interface to the wick-wall interface (superheating) is high enough that boiling can occur within the wick. When this occurs, capillary flow through this point is interrupted, and the portion of the heat pipe deprived of liquid resupply from the wick will dry out. A reduction in heat flux below the critical heat flux is necessary to allow rewetting of the wick and restoration of operation in the heat pipe regime.

The *capillary limit* is generally the most common limiting phenomenon for heat pipes. It occurs when the pressure condition in Equation 2.2 is not satisfied, whether as a result of capillary limit decreasing or one of the other pressure drops increasing. Simply put, where the mass flow of liquid to the evaporator (delivered by capillary pumping) is less than the mass of liquid evaporating from the surface of the wick, dry-out occurs. Slowing the rate of evaporation (i.e. reducing power) is generally the simplest cure for the condition.

Heat Pipe Failure (Irrecoverable)

Unlike the recoverable failure modes, which are the result of limits fairly easily characterized with either basic physics or semi-empirical models, the risk of irrecoverable failures is much more difficult to quantify.

As believed to be the case on the GOES-17 satellite, the presence of particulates in a heat pipe can result in the plugging of wick pores, and the blockage of capillary action. While

there may be some particulates that are sufficiently volatile that they could be forced out of the wick under certain conditions, it must still be considered irrecoverable. So long as the particulates continue to exist in the heat pipe, they will continue to be redeposited in the wick. However, in a heat pipe with excess wick capacity, it may be possible to preferentially deposit particulates where they will not be detrimental, perhaps in an inactive region beyond the primary condenser section. Identifying the presence of particulates on manufacture is possible with electrically-heated testing, provided processes such as brazing, welding, or nuclear effects do not produce particulates later on in assembly or operation.

Impurities that result in corrosion may also be produced as a result of mass diffusion from the heat pipe structure. Reid identifies a number of corrosion mechanisms, and their effect on alkali metal heat pipe life.

In addition to plugging of the wick, it's physical failure, whether from corrosion, rapid acceleration, or other phenomenon, is possible. Its likelihood would increase with complex wick shapes traversing bends in a heat pipe, where unaccommodated thermal expansion might impose stresses on the wick. Wick failure is also generally more probable for brittle wick structures. This failure mode represents a particular concern, however unlikely, as the conditions affecting one heat pipe likely affect all of them. Such failures undermine the fault-tolerance of having a multitude of heat pipes.

Similarly, a breach of the sealed heat pipe volume represents a catastrophic failure. Specifically, any failure that allows the working fluid to escape, in gas or liquid form, or allows the entry of foreign material, permanently prevents operation in the heat pipe regime. Among possible causes of breaches, foreign objects entering the heat exchanger present a significant risk, particularly for systems with direct air cooling and other gas cooled systems with turbomachinery. With high flow rates and sensitivity to pressure drops in the gas, there are limited options for preventing debris entering. Further, with the high flow velocities necessary for efficient heat removal, any debris introduced could impact one or more heat pipes with significant momentum. Likely sources of debris could include external objects (e.g. dust, sand, etc., particularly for a direct air-cooled system) or from the turbomachinery and other components in the flow loop. If heat pipe breaches by foreign material are credible, the possibility of a multiple failure scenario likely is as well.

It is also possible to have a breach as a result of stresses on the heat pipe structure, including those from overpressurization (i.e. from high heat pipe temperatures), stress corrosion cracking, stresses from thermal expansion of various components, and mechanical loads on the heat pipes (such as when a non-stationary reactor is accelerated).

2.3 Principles of Rocket Operation and Evaluation

Rockets operate on the classic momentum balance known from Newton's second law: that the rate of change in momentum of an object is proportional to the force applied. A vehicle, containing a propellant on-board, exhausts that propellant through a nozzle, such that the vehicle moves in one direction and the exhaust in the opposite direction. The flow exiting the nozzle is supersonic, and the basic nozzle parameters, including the exit velocity, v_e , can be determined using the *de Laval* rocket equation.⁹

There are several methods for deriving the classic rocket equation, beginning with a simple description of the momentum balance of the system (i.e. vehicle with decreasing on-board propellant mass, and propellant exhausted with some mass and velocity relative to the vehicle). That equation takes the following form(49):

$$\Delta V = v_e \cdot \ln \frac{M_o}{M_i} \quad (2.10)$$

In this equation, v_e is the velocity of the exhaust gas, M_0 is the initial vehicle mass including propellant, and M is the mass of the vehicle after that propellant is ejected. ΔV is the change in velocity of the vehicle and remaining propellant.

M_0 at any point in time consists of a vehicle or payload mass, plus the total initial mass of propellant. We can redefine M_0 as the sum of a vehicle mass, M_i , and the mass of propellant exhausted, M , and express Equation 2.10 as

$$\Delta V = v_e \cdot \ln \frac{M + M_i}{M_i} \quad (2.11)$$

⁹Note that in this relationship, molar mass appears in the denominator. The maximum exit velocity occurs with minimum molar mass, with hydrogen gas therefore offering the best (theoretical) performance.

and solve for propellant mass needed for a given ΔV :

$$M = M_i \cdot \left(\exp \frac{\Delta V}{v_e} - 1 \right) \quad (2.12)$$

Specific impulse, I_{sp} is commonly used in referencing a given engine's efficiency in utilizing propellant, and is related to exhaust velocity by the gravitational constant:

$$I_{sp} = \frac{v_e}{g_0} \quad (2.13)$$

I_{sp} is typically expressed in seconds, with a larger number representing greater force per unit of propellant.

Specific Impulse on Existing Rocket Engines

For evaluation later of nuclear rocket systems, some comparison to existing chemical rocket systems is necessary. The US Space Shuttle's main engines (RD-25) have the highest vacuum specific impulse of any rocket with significant operating experience at 452 seconds. The reference to *vacuum* specific impulse distinguishes between a rocket which can use its atmosphere as a source of propellant, as is the case with a combustion engine using air as its source of oxygen, and an engine operating in vacuum where its entire propellant mass is carried on-board. Engines not designed for leaving Earth's atmosphere often have much higher specific impulse, particularly at speeds less than Mach 2, than the most efficient thrusters used in space (e.g. turbofans on modern commercial aircraft). Newer, privately-developed rockets such as SpaceX's Raptor and Merlin engines have vacuum specific impulse between 310 and 380 seconds.

2.3.1 Nuclear Rocket Engine System Operation

At the most basic level, a nuclear thermal rocket (NTR), or nuclear thermal propulsion system (NTP) consists of a nuclear reactor which directly heats a gaseous propellant for expulsion through a nozzle producing thrust. Typically, the propellant considered is

hydrogen, stored in a tank on-board the rocket, and pumped through the reactor using a turbopump.

Based on the exit gas temperatures that may be achieved using this technology, performance (in terms of propellant utilization efficiency) on the order of double that of a conventional chemical rocket may be possible.

In most of the rocket engine systems built or examined to-date, there are a few basic features they share. Hydrogen passes directly through the reactor core, heating from less than 200K to greater than 2000K while flowing within small coolant channels in fuel elements having lengths on the order of 1 meter.

In addition to providing thrust, some of the hydrogen is also used to cool certain components, such as the tie-tube structures in the core, reflector and control drum elements, and the engine's nozzle. Typically, the heat from cooling these components is used in an expander cycle to drive the pump that feeds hydrogen into the reactor core. The turbine exhaust on the pump is exhausted into the reactor inlet.

There are numerous other flow configurations that are equally viable, but most of the effort has been on expander cycle engines. For the purposes of reactor system design, the biggest impact of this configuration in a low-enriched design lies in the thermal requirements imposed on the tie-tubes, in this case ensuring enough heat is added to that flow circuit to drive the turbopump.

The tie tube is a structure that serves primarily to hold the fuel elements in place, anchored at the inlet (cold) end of the reactor core, and reaching toward the hot end of the fuel elements. A platform at the hot end of the tie tube engages a notch on the fuel elements to hold them in place. These tie tubes are metal, not graphite, and are unable to withstand the extreme temperatures of the fuel elements directly. So the tie tube is constructed as a set of concentric tubes, with center and annular flow channels. A moderator, also kept at temperatures significantly below the fuel, is placed between the two coolant passages, and heating in the moderator is removed by the flowing hydrogen gas. A basic flow diagram for the Small Engine (a concept from NERVA) is shown in Figure 2.2.(27) Methods used for design and analysis of nuclear rockets are discussed in appendix A.




Figure 2.2: Representative nuclear thermal rocket flow schematic, Small Engine concept.

2.3.2 Nuclear Electric Propulsion

A nuclear electric propulsion (NEP) system is in many ways a much simpler concept. Here, a nuclear reactor generates electricity, and electricity is used to power an ion thruster. Generally speaking, a nuclear electric propulsion system is expected to make more efficient use of propellant, at the expense of lower thrust.

Unlike chemical rocket systems, where specific impulse is generally 300-500s, or nuclear thermal rocket systems, where specific impulse (theoretically) reaches nearly 900s, ion thrusters have operating experience in excess of 4,000s. This translates to an exit velocity of approximately 40 kilometers per second.

In addition to better propellant usage, NEP systems have a very practical advantage over NTP counterparts: the ability to separate development and testing of the reactor and rocket components.

Chapter 3

Reactor Development Plan Targeting Low-Enrichment Nuclear Rockets

Recent advances in materials technology have enabled the design space for nuclear rockets to include low-enrichment fuels. While a highly-enriched system will almost always be lower mass (given extremely low burnup), the advantages to an LEU system are numerous. They are mostly non-technical advantages, including the distancing of the system from weapons applications, enabling commercial shipping and import/export channels, multiple sources of fuel fabrication services, and lower safeguards attractiveness classification. These advantages ultimately point to a lower-cost system, at least in principle. Given that any nuclear rocket system launched will be on a flagship mission, the mass penalty (and cost) associated with an LEU system likely would be dwarfed by costs savings associated with security for vehicle integration alone.

Key in testing any nuclear system is identifying a facility capable of handling its operation. No such facility exists in the United States today. A low-enriched system does, however, significantly expand the number of options for testing, by including large-scale national security installations (e.g. Nevada Test Site) as well as sites suited to rocket testing that currently do not conduct nuclear operations (e.g. NASA-Stennis), without overly burdensome security requirements.

3.1 Development Plan Overview

The remainder of this chapter presents an overview of the proposed testing phases and their purpose. Given the complex, integrated nature of the nuclear and thermal-hydraulic performance of a nuclear rocket system, they are not well suited to electrically-heated non-nuclear testing. Flow rates are strong functions of fuel element temperatures, element power (i.e. localized fission rate) is tightly coupled with flow characteristics. Further, an entire class of fuel (and the only one with any operational history) is composite: fuel kernels are distributed throughout a graphite matrix. This non-uniformity in power generation within the fuel means that inductively heated fuel element testing will be of limited value in determining true failure mechanisms for the fuel. Even with the development of very sophisticated material performance models, the confidence in the design will likely remain too low for launch (especially in the case of a human-rated system) without some form integrated testing here on planet Earth.

Without a history of development and testing, a full-scale test would be very risky: large amounts of money would be spent toward a test, with no useful product if the entire test is not developed and operated.

The plan for arriving at a nuclear rocket test embraces the incremental testing approach started with the DUFF experiment. It advances through a series of alternating new nuclear system tests and component-level tests, culminating in a nuclear propulsion system test, operating in a closed loop in the case of a thermal rocket. The design of the closed-loop test will be such that any fuel fragmentation or failure will be contained within the facility.

The objective at each phase of testing is to address a manageable set of programmatic and technical risks, while operating several near-prototypic power systems along the way. This serves several purposes: (1) key resources are spread across a limited number of development and planning activities, (2) regulators see a clear connection and incremental step beyond a previously operated experiment, and (3) facilities and test articles useful beyond nuclear rocket applications are put into operation, garnering support for the effort from outside the space nuclear propulsion community. Further, should there be an interruption or cancellation of program funding, the completed facilities and tests still contribute useful knowledge,

through the reduction of theory and modeling to practice. Too many prior space nuclear system design efforts have been shelved in paper-only form, which leaves behind only the rationale of what should work, not what actually does.

3.2 Sub-Kilowatt Nuclear Demonstration

The first phase of the development plan is to do something, *anything*, testing a nuclear system. The goal is to minimize both the technical and political risk to get over the threshold of having operated a nuclear system at all. This phase of testing has been completed, and both its development and the execution of the plan are detailed in chapter 8.

3.3 Sub-Kilowatt Non-Nuclear and Nuclear Tests

The next phase of testing will consist of component-level electrically-heated testing of the concept associated with the DUFF test, the Simple Fission Reactor. This is a compact, $500W_e$ system utilizing a solid metal core of highly-enriched uranium surrounded by a beryllium oxide reflector containing a set of sodium heat pipes. In the concept, these heat pipes transport heat to Stirling engines (8), rejecting heat to space via NASA-designed radiator panels.

The test will consist of a depleted uranium fuel block, manufactured at Y-12, reflector, heat pipes, and at least one Stirling engine pair. It is operated at prototypic temperatures (1100K Stirling hot-head temperature), and completely in vacuum.

Early planning of this test was undertaken in late-2012, with a team consisting of the author supported by staff members at NASA-GRC, DOE Y-12, and Los Alamos National Laboratory. Like the DUFF experiment, the nuclear operation of this test would occur at the Device Assembly Facility (DAF) in Nevada. Beyond that, however, the author's approach differs from what would become KRUSTY.

It would make use of a general-purpose lift table, and its associated control and safety systems, to raise and lower a portion (the lower section) of the reflector. The control rod

intended for a flight system would be omitted. This system would ideally also be conducted with all but the movable reflector section inside a vacuum chamber.

It would demonstrate the full heat transfer mechanism of the Simple Fission Reactor concept, with heat pipes external to the fuel in the radial reflector. It would, as a result, be very near flight-like, replacing only the lower reflector with control rod and drive. If desired, the reflector segments could be machined so that the exchange of those two components could be a bolt-on operation. If for some reason the political will existed, and the hardware was flight qualified, it could be identical to a unit delivered to a launch pad.

This is a significant step forward from DUFF, requiring the approvals to operate an entirely new reactor system at a much higher continuous thermal power than any other experiments in the facility (DUFF was previously the highest), and requiring modifications to the authorization basis for the facility. As it existed for DUFF, an \$0.80 cold excess reactivity limit was in place on the non-pulsing experiments. This is primarily for safety reasons, in administratively ensuring that prompt criticality is precluded. Like with DUFF, excess reactivity is measured/calculated at the start of the experiment to confirm consistency with expectations. However, the large negative thermal feedback that must be overcome to get to the elevated temperatures needed for the test will require more reactivity. This was the kind of modification that should be largely a matter of routine activity, but any such required approvals always carry programmatic risk.

In addition to the new reactor system, this experiment also tests the procedures and obstacles to getting beryllium-containing parts into the facility. Like DUFF, it is intended as a temporary installation and setup, operated through a series of tests, and dismantled to return the lift table to normal service. Since the author's involvement concluded, it grew into the KRUSTY/Kilopower testing recently completed at the Nevada National Security Site.

It has since grown in scope, increasing in power to a $1kW_e$ design, which increased its initial estimated cost to approximately \$20 million, and with that the test article has deviated significantly from a flight-like prototype. As completed, the effort was more expensive and took significantly longer than what the author advocated for. However, the larger-scope experiment nevertheless incorporated most of the intended objectives outlined above.

The system tested by NASA and LANL differs from the one here. Retreating back to the $500W_e$ power level would have allowed the test to more closely mimic a flight-prototypic system in several key areas. First, its increased power and size of KRUSTY resulted in relocating the heat pipes from the peripheral reflector to the exterior of the fuel itself. Second, there was no mechanical connection between the reactor core (i.e. fuel and heat pipes) and the reflector. Given the significant neutronic worth of the reflector, particularly as it moves with thermal expansion, it would have been valuable to include this phenomenon in the demonstration. Third, only part of the system was operated inside a vacuum chamber.

Advancing the Simple Fission Reactor to flight status would benefit from further testing; however, it no longer offers programmatic value in working toward larger space nuclear power systems or portable ground-based reactor systems.

The concept is important, however, in advancing toward a power system actually operating in space. Viewed in light of mission planning, and the ever-changing big-picture plans for the nation's space agency, getting a simple low-power system into space as an alternative to other power sources, is more important than the next return-to-the-Moon or manned Mars mission. Creeping toward larger and more complex concepts is also creeping out of near-term launch plans.

3.4 Phase II: Fission Surface Power (FSP)-Derived Test

This test sets out to develop a reactor in a facility where it could be operated for extended periods of time at a reasonable cost, while still operating at temperatures consistent with the Small Fission Reactor concept and KRUSTY demonstration. The rationale here is to continue advancing and proving real nuclear power systems, with applications in space, while providing useful life testing data that would be applicable to operations associated with (and flight qualification of) a deployed Kilopower system. The reactor would take the form of a heatpipe-cooled system utilizing Stirling engine power conversion, like DUFF, KRUSTY, and Kilopower, but its operation at higher thermal power ($10kW_t$ to $200kW_t$) would drive the

design to discrete fuel elements. They would be installed in cylindrical form in a hexagonal array alternating between fuel and heat pipes.¹ The heat pipes would be integral to the block, constructed of stainless steel. Control of the reactor would be via rotation of approximately 6 external control drums, composed of beryllium oxide and boron carbide, with a 2-of-6 operation requirement for shutdown. It would also contain a central cavity for a shutdown rod.

The key achievements of this test will be the fabrication of an integrated core-heatpipe block, extended operation, introduction of a new control system (and authorization basis) and the first domestic reactor in many decades to utilize rotating control drums. In the process, it would be providing operational experience and life testing data relevant to a surface power application as would be needed for extended stays on the moon or Mars, and provide a fast reactor research facility otherwise not available domestically. If constructed and operated in a new location, it would also exercise the licensing process for research reactors, which would prove beneficial in approaching the large-scale tests later.

Given that such a system could be designed with a lower safeguards and hazard rating than what exists at the DAF, a less expensive (long-term) location is necessary. Existing facilities used for the previous experiments are not sufficiently available for long-term use, from either a practical (access, radiation exposure, etc.) or programmatic (cost, displacement of other basic science and national security research) standpoint. This would include possible operation on or near a university campus where it could also support other fast-flux research. Selecting a facility where it could be operated for extended periods of time at a reasonable cost, while still operating at temperatures consistent with the Kilopower prototype is key to its longer-term success. In terms of the broader propulsion objective, this system sets the stage for the first new modern reactor facility with aspirations for operation over extended intervals of time.

¹Alternatively, if demonstrated separately early enough in the project, the system could make use of one or more heat pipe configurations described in Chapter 5.

3.4.1 Facility Needs, Licensing Process, Estimated Facility Costs

Proposing a new research reactor licensing effort carries enormous risk, especially given the time that has elapsed since any new reactor has gone through the process. A class 104(c) license, for research and development facilities, is the simplest licensing process for a reactor system of this type. The power produced, even if operated continuously for life testing, is of minimal value; the costs of a part 103 license, even as a prototype, would exceed the value of the energy sold. While 104 is an *easier* licensing process than part 103, it is not *easy*: the description of the format and content of an application exceeds 500 pages.

Since the concept proposed is substantially different than any NRC-licensed reactor (commercial or research and development), and the process has not been exercised in some time for a research reactor, the proposal here takes a stepwise approach to licensing. The first stage in licensing is a foot-in-the-door step: licensing a new, zero-power critical assembly. As knowledge and experience are gained, the licensing basis is amended to address longer duration operation at temperatures, increase excess reactivity limits, and to begin operation with the power conversion system. Construction of a system of this type should be able to be licensed for less than \$5 million, and constructed for less than \$20 million.

3.4.2 Conceptual Design

As a reactor grows in size, the monolithic block fuel forms used in the Simple Fission Reactor, DUFF, and kilowatt-class reactors like KRUSTY, fail to satisfy operational requirements. Total fuel exposure (i.e. burnup) over a many-year lifetime is at levels where fuel swelling and fission gas release are significant concerns. Core linear heat rates (i.e. total core power per unit height) are sufficiently high that cooling only from the periphery would result in unacceptable peak fuel temperatures. Introducing heat pipes within the fuel creates a number of fabrication challenges, mostly with regard to how a heat pipe is inserted into the fuel block with sufficient (and sufficiently reliable) thermal contact.

3.4.3 Design Process

As discussed in section 2.1.2, thermal systems (particularly undermoderated thermal systems) are very difficult to design such that they satisfy criticality requirements for launch accidents. If minimizing fuel mass, particularly if enrichment constraints exist, is a primary design driver, a UZrH-fueled system can work well in this power range. Over longer lifetimes, hydrogen migration can result in a loss of reactivity, shortening system life.

The key achievements of this test will be the fabrication of an integrated core-heatpipe block, extended operation, introduction of a new control system (and authorization basis) and the first domestic reactor in decades to utilize rotating control drums. In the process, it would be providing operational experience relevant to a surface power application as would be needed for extended stays on the moon or Mars, and provide a fast reactor research facility otherwise not available domestically. If constructed and operated in a new location, it would also exercise the licensing process for research reactors, which would prove beneficial in approaching the large-scale tests later. The performance of this system can be improved significantly using the multi-stage heat pipes in Chapter 5.

3.5 Megawatt-Class Terrestrial Reactor Test

To help sustain the program as political interests in space applications, particularly of the scale that would demand a nuclear propulsion system, prove weak during this process, an alternate (possibly parallel) path to the nuclear propulsion test can be undertaken addressing terrestrial needs. Here, a larger heat pipe reactor, using similar construction methods as the FSP concept, would be built to produce approximately 5MW_t . It would be targeting manned installations (e.g. forward operating bases for the Army) in remote locations, replacing large portable diesel generators and/or so-called “micro-grid” systems. Depending on the political climate at the time, these systems may also prove attractive for more conventional grid-tied power generation as well.

Key differences between this and the previous systems include breaking the reactor into 6 wedge-shaped assemblies and utilizing low-enriched fuel, possibly including high-assay low-enriched uranium. Each assembly would contain fuel, heat pipes, control and reflector,

heat exchangers (depending on power conversion type), power conversion and heat rejection that would be independent of the other assemblies, but would require the other assemblies to approach a critical system. This would allow transport of the assemblies separately with minimal criticality safety concern prior to operation, while maintaining a very short installation time/effort. In the process of designing this system, the innovation that now stands alone as Chapter 5 was developed.

3.6 Evaluation of Nuclear Thermal or Nuclear Electric for Downselect

Until this point in the testing phase, all of the reactor system development has had multiple purposes. The reactors all address technical and programmatic needs for the nuclear thermal propulsion interest, but have usefulness well beyond those applications. However, going forward with a large-scale closed-loop nuclear propulsion test will have limited appeal beyond space propulsion, and only very limited relevance beyond similarly designed (e.g. fuel form and engine cycle) thermal propulsion systems.

Given that even full implementation of this plan without delays would take on the order of 20 years to this point, and that electric propulsion technology continues to advance with radioisotope and solar power as the energy sources, a nuclear electric propulsion system is likely far more useful. Depending on advances in that technology, and advances in high temperature materials that could be used in building a lightweight high-power reactor, it is quite possible that advancing the megawatt-class system proposed to a low-mass megawatt-class system is a routine engineering effort.

As research into thermal propulsion systems has continued, the author has become even more convinced that a large-scale NTP test would not be justified. In some respects, the development plan, in hopes of an NTP test, advances the cause of space exploration while decreasing the likelihood of NTP deployment. It does so by providing the facilities to develop the knowledge and experience for a competing system.

3.7 Large-Scale Nuclear Propulsion Test

This system was intended to be the primary design focus, making use of the improved nuclear rocket design process described in the appendices.

Given the marginal improvements over traditional rockets, significant reductions in Earth-to-orbit costs as a result of commercial launch capabilities, and inferior performance relative to electric propulsion, a nuclear thermal rocket system as conceptualized since the early ROVER program makes little sense.

In addition to the rocket design process, since exhausting hot hydrogen that may contain fuel fragments will be unacceptable here on Earth, testing is also complicated. It will require a large-scale hydrogen capture system to prevent radiation releases to the atmosphere. Fortunately, nuclear rocket systems do not operate for long periods of time, so the capture system does not have to maintain a steady-state combustion, cooling, and condensation process. Even so, the capture system will be much larger than the reactor system itself.

Recent analysis of nuclear thermal reactor systems' performance in postulated launch accident scenarios by Poston ((38)) suggests that preventing criticality in these accidents will require complex systems and operations, such as withdrawing bundles of fine wires inserted into the fuel element cooling channels. When combining the complex, untested/unqualified fuel forms necessary to meet required operating temperatures, systems required to prevent criticality on launch, along with all of the usual complexities of engineering a rocket engine, a factor-of-2 performance gain isn't sufficient to justify the billions of dollars necessary to deploy such a system.

Chapter 4

Megawatt-Class Terrestrial Reactor Demonstration

Since the Demonstration Using Flattop Fission (DUFF) experiment, the first heatpipe-cooled nuclear reactor power system, was operated in 2012, interest in practical Earth-bound applications of heat pipe technology have grown. The success of the two heat pipe reactor systems operated to-date has renewed interest in megawatt-class, portable heat pipe reactor concepts, particularly in the context of remote locations, supporting forward military needs, and as part of disaster response efforts.

This phase of testing intends to advance the state of the art in development of this class of portable, compact reactor systems with a design configuration that enhances modularity, and demonstrates its feasibility, first through non-nuclear testing using state-of-the-art turbomachinery, followed by full-scale deployment of a test reactor. Unlike the small heat pipe reactor systems that have been previously operated, the large number of heat pipes in a megawatt-class reactor system present certain unique challenges.

Heat pipe reactors, here again using alkali metal heat pipes, are perfectly suited for mobile applications because their design is inherently simpler, smaller, and more reliable than “traditional” reactors that rely on pumped coolant through the core. Despite acknowledged benefits associated with heat pipe reactor systems in general, concerns have been expressed with regard to the possible use of heat pipe reactors for applications at powers greater

than 1MW , namely the possibility of cascading failures and the mechanical complexity of integrating a large number of heat pipes to a dynamic power conversion system.

As heat pipe reactor concepts grow in size, many include multiple core block sections, separately manufactured, connected to a common heat exchanger or plenum for connection to a single power conversion system. The smaller individual block sections are generally easier to manufacture, and with fewer machining and welding operations per block, the consequences of machining or welding defects are reduced. But in such configurations, the complexity of connection to a common heat exchanger remains. The use of additive manufacturing for some of these parts may mitigate some of these problems, but blocks containing fewer fuel element and heat pipe penetrations and heat exchanger connections will likely remain more favorable.

One solution to the mechanical complexity of integrating a large number of heat pipes with a dynamic power conversion system is to instead integrate a smaller subset of heat pipes with independent power conversion systems, as shown in Figure 4.1.

Heat pipe reactor systems are not inherently linked to a particular power conversion type; however, concepts targeting rapid, remote deployment and/or micro-grid operation generally rely on ultimate heat rejection to air. Direct utilization of air in an open Brayton cycle is not the most efficient embodiment of a heat pipe reactor, but reduces the system to one where only the fluid in the sealed heat pipes is necessary for system operation. A more efficient system could be constructed with a closed-loop Brayton system, rejecting waste heat through more conventional radiators.

In the context of more conventional reactor systems, and even the majority of near-term advanced reactor concepts, this design approach eliminates one of the most significant operational risks with a nuclear power system, namely the loss of primary coolant and/or the failure of the ultimate heat sink. With this concept, the direct utilization of air means that so long as sufficient heat is produced, the turbomachinery will continue to move air across the heat pipes and cool the core, even in the absence of control or external power. This is in addition to a passive decay heat removal system that directly rejects heat to air with natural convection.




Figure 4.1: Segmented core concept (top), with individual core modules connected to each other via a common heat exchanger (containing many more brazed/welded joints) as incorporated in (30). Proposed concept (above) has only electrical interconnections between modules, allowing for more rapid deployment and transport in smaller packages if necessary. This approach also reduces the number of heat pipe integration operations per core block.

A heat pipe reactor sub-module containing a core block, with fuel and heat pipes extending out of the active core, connecting to small air Brayton systems and alternators, would comprise one of several identical subsystems that could be linked together to operate as a larger-scale reactor system. With only one sub-module, the system is incomplete and does not present a criticality safety risk. When grouped as a cluster, and surrounded by a reflector and control system, the modules can operate as a complete system. Because each sub-module would only need power and instrumentation connections to the rest of the system when set into place, the reactor could be transported into areas where a heavy truckload is not feasible, with minimal assembly effort onsite. In other words, a truly modular reactor system—composed of standardized units for easy assembly, as opposed to a larger pre-assembled reactor system.

The concept is novel, and takes advantage of advances in turbomachinery over the last decade that have significantly improved operating characteristics to the point that commonly available automotive examples are now capable of higher temperatures than a heatpipe reactor system would need. What once might have been cost prohibitive for a small reactor, .e. having six or twelve sets of superalloy turbomachinery instead of one, is perhaps now an advantage.

The earlier intention with this section was to simply develop a concept far enough to get to a testing plan. As it developed however, a novel heat pipe and fuel design emerged. Analysis of this element is presented in Chapter 5.

4.1 Facility Needs, Licensing Process, Estimated Facility Costs

Unlike the previous tests, where the power produced is relatively small, the megawatt-class system is primarily designed to supply electricity to an electrical grid or micro-grid. As a domestic system, using it in this form is no longer permitted by a 104 license, and the usual commercial reactor licensing processes apply.

A part 103 license is significantly more involved than part 104(c), and there are numerous efforts underway by potential small modular reactor operators and designers to reduce the burden of a commercial license. Several (including Oklo and mPower) have attempted to partner with the Department of Energy to make use of sites already familiar with nuclear operations (including Oak Ridge National Laboratory and Idaho National Laboratory) where land use is already regulated and the impact to members of the general public is limited.

Unfortunately, this is a licensing process not well exercised in recent decades. Even attempts to build a largely conventional reactor in the last decade have failed after investing billions of dollars, and those currently under construction are significantly behind schedule and over budget. There are a variety of causes, from supply chain issues to changes as a result of the earthquake affecting the Fukushima Daiichi nuclear power plant in 2011.

The plan here proposes construction at or adjacent to an existing commercial nuclear plant site, not at a national laboratory, for two specific reasons: (1) the siting requirements, including emergency planning zones, have already been evaluated by the NRC, and (2) exercising the usual licensing processes, including those requiring public input, is necessary to prove the system worthy of investment and wider adoption. Since the licensing process will span multiple years, that new buildings will need to be constructed, and that regulated utilities will be involved, this is a significant undertaking. It should be expected that the cost to first-of-a-kind operation will cost at least \$100M.

4.2 Scaled Non-Nuclear Testing

4.2.1 Required Facilities for Non-Nuclear Testing

Existing design tools and computing resources are available to advance the system concept design, both in terms of the nuclear system (using MCNP and supporting design automation tools) and the power conversion system (with existing commercial computer-aided design software). The test apparatus can be manufactured with existing conventional and CNC machining capabilities along with additive manufacturing systems now common to many lab and manufacturing spaces.

Assembly and operation of the test will require lab space, power, and an ability to reject the heat generated by the system. Since it is an air-sourced, air-rejected power system, a small lab space internal to a larger facility would not be adequate. It will need access to fresh air. While this eliminates many facilities as test site candidates, there are many spaces currently vacant with sufficient power available and the ability to open to the outdoors—what’s needed here is not an advanced facility, but instead underutilized manufacturing space. The reactor concept will be operable under a wide range of weather conditions, however, and there may be some desire to operate the system at more than one site.

Chapter 5

Double-Barrier Fuel and Annular Heat Pipe Assembly

5.1 Introduction

This section introduces what the author has termed a *double-barrier fuel and annular heat pipe assembly*. The nomenclature refers to the complete enclosure of the fuel element within a heat pipe, such that both the fuel cladding and the heat pipe wall would have to be breached for any release of fission gases to occur, and to the arrangement of at least the evaporator section of the heat pipe in an annulus *around* the fuel element.

It is presented in several variations, arising out of a desire to shrink the footprint of a typical heat pipe reactor system. To understand the development and its merits, one must understand limitations and complications inherent in existing designs. Let there be no doubt, however, that even heat pipe reactor designs making use of conventional heat pipes represent significant advances in simplicity and reliability, including DUFF and KRUSTY with operational history and concepts such as HOMER(37), the Simple Fission Reactor, and the Mobile Heat Pipe Cooled Fast Reactor System.(30)

5.1.1 Existing Heat Pipe Reactor Concepts

At the smallest scale, and as demonstrated with DUFF and KRUSTY, particularly for space applications where minimizing mass is a key objective, the system is criticality limited. Safety criteria for launch, including subcriticality on immersion in water and wet sand, tend to favor fast-spectrum systems with high reflector worth. The reasoning stands that an accident where the reflector remains fully intact but a central control rod is ejected is not credible, and that any other combination of materials surrounding the core, even of infinite dimension, are insufficient to make that core critical.

Past systems have proposed a solid cylindrical fuel block, with a number of cylindrical heat pipes dispersed within it. Many of the concepts never reached a level of maturity where manufacturing techniques were identified or proven, but most assumed some sort of permanent structural bond between the heat pipe, manufactured and filled in advance, and the fuel block. The heat pipes are responsible for transporting heat generated by the reactor to either a heat exchanger or power conversion system (e.g. Stirling converters).

As power levels (and total burnup) increase, a number of concerns arise, including criticality during assembly, heat transfer to and from the heat pipe surface, means of heat removal in the case of a heat pipe failure, manufacturability, power distribution, and others. To address these concerns, larger heat pipe reactors move to discrete fuel elements interspersed in an array of heat pipes. Both heat pipe and fuel element are cylindrical, with the heat pipes extending out of the core section to pass through shielding and integrate with power conversion and decay heat removal systems.

The heat pipes are each mated by some form of mechanical connection (e.g. brazing, welding) to a monolithic support block that provides the primary mechanism for heat transport between fuel and heat pipe. Fuel element locations are constructed similarly, though the author has previously proposed manufacturing techniques that would use the block as both part of the heat pipe wall and fuel element cladding, a technique that is the baseline for several concepts under development and license.³⁰

In these systems, the pitch-to-diameter ratios and fuel-pin to heat pipe ratios are both constrained by the need to maintain reasonably uniform azimuthal heat fluxes into the heat

pipes, and to provide a conduction pathway to a next-nearest heat pipe in the case of a heat pipe failure. Temperature gradients in the fuel block create stresses, and a minimum web thickness becomes necessary to maintain core block integrity.

Similarly, for a cylindrical heat pipe with a limiting heat flux, the ratio of wick surface area (L_w) to vapor flow area (A_w) is also fixed, and inversely proportional to the heat pipe radius (R_{HP}),

$$\frac{L_{wick}}{A_{flow}} = C \quad (5.1)$$

with minor deviations based on wick type and thickness.

It is important to note that in these conventional designs the heat flux into the heat pipe is not uniform with azimuth, with significant dependence on core block and interface conductance. This is the result of the variable distance between the heat source (fuel pin) and the heat pipe wall. Under normal operational conditions, azimuthal flow within the wick and variable evaporation rates make the effect inconsequential, but it can significantly reduce margin to a given heat pipe's limits when operated near failed or inoperable heat pipes (i.e. removing additional energy delivered from only one direction).

5.2 Impact on Core Layout/Neutronic Design

The simplest form of this new assembly moves the heat source—nuclear fuel—to the center, surrounded by wick and then vapor flow space, contained by an outer wall (ignoring, for now, the arrangement of such a heat pipe beyond the in-core evaporator section). That change adds a degree of freedom in the relationship between azimuthal wick length and vapor flow area, without changing wick thickness. But it offers significant advantages even if we conserve flow area for a given heat pipe.

Figure 5.1 shows a conventional hexagonal region within a heat pipe reactor, with 2 fuel elements for every heat pipe. For each whole heat pipe, 120-degree segments of the neighboring six fuel elements fall within that region. A smaller region forms a unit cell,¹ A , containing one-sixth of a heat pipe, and one-third of a complete fuel element.

¹This is not the smallest symmetric geometry, but is chosen so that the comparison to the novel configuration is clearer.




Figure 5.1: Conventional Heat Pipe and Fuel Layout with Discrete Fuel and Structural Block/Monolith and 2:1 Fuel Pin to Heat Pipe Ratio



Figure 5.2: Combined Heat Pipe and Fuel Unit Cell Equivalent

If one conserves the cross-sectional area of each fuel element and heat pipe, the equivalent unit cell dimensions grow by $\sqrt{2}$, and unit cell area doubles. In that larger space, it incorporates 50% more fuel area and triple the heat pipe area and azimuthal wick length. Recalling limitations on heat flux into a heat pipe evaporator and maximum power density for a given vapor space, this would imply a 50% increase in power density, a 25% reduction in fuel volume, and a 25% reduction in linear heat rate. Similarly, conservation of heat pipe flow area in a given unit cell results in a unit cell enlarged in area by 1/3, containing 50% more fuel, and with three times the wick area. Of course, given the annular geometry, the hydraulic diameter for a given flow area is reduced, and this must be taken into account when making the comparison.

While either of those configurations is interesting, it is more significant that between those two configurations is an infinitely variable parameter space. Heat pipe reactor concepts with discrete fuel elements thus far have adjusted the balance between fuel loading and heat transport by varying the number of fuel elements per heat pipe. With heat pipes removed from the fuel element, and heat conduction through an interstitial block to a heat pipe, decreasing the heat pipe size increases the effective distance between fuel pin and heat pipe, resulting in greater temperature gradients. It also reduces flow area and surface area together. If one aspect of heat pipe design (e.g. capillary or sonic limits) requires an increase in size, the other aspect may become over-designed.

5.2.1 Redundant Barrier Between Fission Products and Atmosphere

The arrangement of a fuel element internal to an annular heat pipe section places a second barrier between the radioactive inventory in the fuel element and the environment. That is to say that a fuel cladding must fail *and* the outer heat pipe wall must fail for any radioactive release to occur.

In the event of a failure of the inner barrier (i.e. the fuel pin cladding), non-condensable gases would be driven to condenser end of heat pipe. If temperature (and pressure with it) is high enough, heat removal would continue, with some portion of the condenser end

“disabled” by the fission gases. It would essentially act as variable conductance heat pipe. This introduces several opportunities to detect a cladding failure prior to any release that might affect reactor operations, or pose a threat to health or safety.

5.2.2 Reduction of Likelihood of Cascading Failures

One of the bigger criticisms of heat pipe reactor systems involves the possibility of cascading failures. The concern is that, should a heat pipe fail to remove heat for any reason, energy from adjacent fuel pins must be removed by the neighboring fuel pipes. Those heat pipes then operate at higher power levels, increasing the likelihood that they too will fail.

Figure 5.3 shows a diagram of a conventional 2:1 layout surrounding a fuel element. Under normal operation, heat from the fuel marked at the center of the diagram is distributed amongst the 3 adjacent heat pipes, contributing $1/6$ of the power of each of those pipes. If one of those heat pipes ceases operation, energy from the adjacent fuel elements must be distributed amongst the 2 remaining adjacent heat pipes, or a 50% increase in the amount of power it contributes to each of those heat pipes. Since there are 2 adjacent fuel elements experiencing this condition, the power delivered to each of the adjacent cells is increased by $1/6$.

Second-order effects mitigate some of this increase. Generally, the heat pipe would be designed to operate in a regime where an increase in temperature also results in an increase in maximum power (i.e. capillary, sonic, or entrainment as limiting phenomenon, not boiling). As the neighboring clusters rise in temperature in response to the failed heat pipe, they also conduct some of the additional energy to neighboring clusters (this conduction, and the associated stresses generated from the elevated temperature gradient, are the basis of the design-limiting stress condition for the structural block).

Initially, it would seem that little changes in the annular heat pipe configuration, as shown in Figure 5.4. A failed element must reject all of its heat to six adjacent elements, increasing the power delivered to each by $1/6$. However, that power is being deposited in the outside wall of the adjacent heat pipes, conducted through a block which under normal operation sees no heat flux. The vapor flow rate is increased by $1/6$, pushing closer to the heat pipe’s sonic limit (generally, normal operation is in a regime where the sonic limit is




Figure 5.3: Conventional Heat Pipe Reactor Failure showing Power Delivery to Adjacent Heat Pipes

far in excess of capillary and boiling limits), while essentially leaving the margin to the capillary limit unchanged. The outer wall, of course, has a greater surface area than the inner wall, and is exposed to a heat flux less than $1/6$ that of the inner wall. (As the design is discussed further, certain additional features will cause deviation from this arrangement while preserving the overall operation.)

In the conventional heat pipe reactor arrangement, the peak fuel temperature is located between the point nearest the failed heat pipe and the center of the fuel element. Location of the peak temperature along that line is a function of web thickness, block conductivity, fuel conductivity, and linear heat rate. Detailed analysis of heat pipe failure in this context has been analyzed in [36](#), and for a $5MW_t$ concept containing 2112 fuel elements, estimated a fuel temperature rise on the order of 40K for an average-power pin. This is in addition to similar gradients between fuel and heat pipe under normal operation, implying a temperature difference of nearly 100K between peak fuel temperature in the failed case and heat pipe operating temperature under normal operation.




Figure 5.4: Annular Heat Pipe and Fuel Unit Failure showing Power Delivery to Adjacent Heat Pipes

In the fuel-in-heat-pipe design, we begin with a much lower temperature difference between fuel centerline and heat pipe operating temperature under normal operation. Since the heat pipe is in direct contact with the outer diameter of the fuel element, the nominally 40K temperature drop through the block is eliminated. A failed heat pipe, however, means there exists a gap between the outer surface of the fuel element and the inner surface of the heat pipe/block. Particularly for space applications, such a gap would have very poor thermal properties if the heat pipe fails due to a breach, especially with a freestanding wick structure impeding radiative transfer.

Recalling prior discussion of *failed* versus *inoperable* heat pipes in 2.2.1, the design of a heat pipe is constrained primarily by heat transfer limitations where the working fluid remains in the heat pipe. Outside of the smallest reactor systems designed to operate for many years with no reactivity control input, the response to an inoperable condition is to reduce power and re-wet the evaporator portion of the wick.

For the more extreme conditions where fluid is permanently lost, or where operational constraints preclude rewetting the wick, a conduction pathway is necessary. Given the scale of the conduction pathway (i.e. the radial thickness of the vapor space), however, even a

small amount of contact between the inner and outer surfaces could be enough to mitigate a failed heat pipe. One method of doing so is to break the annulus into segments, with ribs connecting the inner and outer surfaces. For every rib added, flow area decreases, and the annulus thickness must be adjusted.

Assuming, for now, that a design exists for conduction across the annulus in the event of a failed heat pipe, a path forward to a reactor design flows relatively quickly. For a fast spectrum system of a conventional heat pipe arrangement, it is also possible to approximate the equivalent design going forward.

Chapter 6

Figure-Eight Flow Path with Constant Wick Diameter

Advancing the fuel-in-heat-pipe arrangement just described in its simplest form, one would have a cylindrical wick of constant diameter surrounding the fuel element in-core, and located just inside of the outer wall (surrounding the vapor flow space) in the adiabatic and condenser regions. The wick continues beyond the fueled region, through to the condenser region without any change in diameter necessary.

The figure-eight description follows from the crossover that must occur between the wick (liquid) and vapor flow paths, as shown in Figure 6.1. This crossover region begins just beyond the evaporator section, and consists of longitudinal openings in the wick to allow vapor to pass from the outer flow region to the inner flow region, as shown in Figure 6.2.

Beyond the crossover region, the heat pipe operates as any conventional cylindrical heat pipe would, with vapor condensing on the inside of the heat pipe wall and the resulting liquid transported back toward the fuel element. The return path in the wick encounters a reduction in area in the crossover region, which reduces the maximum capillary pressure.

6.1 Design Constraints

In the crossover region, there are several competing design constraints. First, the cross-sectional wick area is reduced by the area opened for vapor crossover. Second, the vapor




Figure 6.1: Flow diagram showing figure eight heat pipe. Not to scale; ratio of heat pipe length to diameter much greater than shown.




Figure 6.2: Diagram showing crossover region of figure eight heat pipe.

flow area in the axial direction is doubled. This reduces the pressure drop in the vapor flow path, recovering margin to the capillary limit (though generally negligible under normal operating circumstances).

$$\Delta P_{vapor} = \rho_v \cdot v^2 \cdot L \quad (6.1)$$

$$v = \frac{\dot{m}}{A_v \cdot \rho_v} \quad (6.2)$$

One quickly sees that in doubling the flow area, pressure drop in the vapor is reduced by a factor of four. Over the same length, we have a reduction wick area, proportionally reducing the maximum power that can be pumped by the wick (see equation 2.9). Given that boiling within the wick is not a concern in this section, one solution is to add a secondary, high-permeability wick on the inside, reaching from the condenser to just below the crossover region.

This requires the crossover region to be designed with two constraints: (1) that the reduction in cross-sectional wick area is small enough that the capillary limit is not exceeded, and (2) that the length of the wick openings is large enough that the pressure drop in the radial direction (through the wick openings) is negligible. Increasing the length of the crossover region further reduces the pressure drop and increases the maximum power the heat pipe can carry.

6.2 Physical Assembly and Preventing Entrainment in Crossover Region

Fabrication of this fuel-in-heat-pipe assembly is no doubt more complex than assembling a standard cylindrical heat pipe. That said, the additional complexity does not dramatically increase the number of components or assembly steps, and may result in more predictable manufacturing losses.

The wick as previously discussed is a simple cylindrical wick, deviating only in the introduction of longitudinal openings in the crossover region. The outer wall is constructed

in two pieces, first the condenser tube, which is just larger than the wick diameter, and extends from the condenser end to the top of the fueled region, with openings matching the openings in the wick. By "covering" the wick in the transition area, interfacial shear stresses are avoided, and the tabs are bent into place, attaching to the end of the fuel cladding, helping to locate the upper end of the fuel element in the center of the heat pipe.

A second tube, fitting inside the wick, makes up the fuel clad. The upper end is closed, while the bottom end remains open for fueling after the heat pipe is filled and sealed. A third tube, matching the diameter of the heat pipe in the evaporator (core) region, having a chamfered upper end for welding to the condenser tube, is inserted over the bottom of the fuel element. The heat pipe is then purged and filled, and its operation tested before any fuel is inserted.

6.2.1 Drawbacks

There are two concerns with this configuration. First, and already mentioned, is the need for a conduction path between the fuel element and the heat pipe periphery, permitting heat rejection to adjacent assemblies. Second, while the cylindrical wick is more than capable of transporting liquid back to the evaporator region, condensation may occur on the inside of the outer heat pipe wall without a mechanism to transport the liquid back to the inner heat pipe (outer fuel element) wall.

Both drawbacks can be addressed with a single element, in the form of radial ribs connecting the inner and outer walls. Figure 6.4 shows one form of those ribs, which facilitate a minor capillary flow from the outer wall to the inner wall. The rib structure adds a conduction path between the inner and outer radii in the absence of the working fluid, while the grooves promote fluid return in the event of condensation on the outer wall (which could occur in situations where the outer wall is cooler than the inner wall).

Further improvements are still possible. An additional wick could be installed on the inside of the outer wall, but would only reach so far as the crossover region. Extending the larger diameter the full length of the heat pipe results in an annular heat pipe configuration, which is introduced in the next section.




Figure 6.3: Flow diagram showing straight annular heat pipe. Diametrical wicks shown in cross-sectional view. Not to scale; ratio of heat pipe length to diameter much greater than shown.

6.3 Straight Concentric Annular Heat Pipe with Diametrical Wicks

This form of heat pipe evolves from the figure-eight layout in the previous section, but continues the larger diameter along the full length of the heat pipe. Wicks on both the inner and outer walls increase the pumping power (for a given wick thickness), and short wick structures provide both a conduction pathway for heat removal in a breached heat pipe and a liquid flow path from the condensing surface on the outside of the heat pipe to the evaporator surface on the inner wall of the heat pipe. One similar structure has been studied for use on laser printer fusers, where rapid heating and uniform temperatures are required⁽²⁴⁾⁽²⁾, along with others where heating and cooling occur on both the inside and outside surfaces.⁽⁵⁰⁾ The earliest known annular heat pipe geometry was patented by Kirkpatrick in 1972.

Compared to a conventional heat pipe, the additional cross-sectional wick area doubles the maximum power that can be transported without reaching the capillary limit (recalling equation 2.8).¹ Wicking between inner and outer walls via diametrical wicks (see Figure 6.4) incurs negligible pressure drop in the liquid phase due to the short path length, but does

¹This is assuming a design where wick thickness is constrained by the boiling limit.




Figure 6.4: Flow diagram showing straight annular heat pipe. Diametrical wicks shown in cross-sectional view. Not to scale; ratio of heat pipe length to diameter much greater than shown.

reduce the area available for vapor flow. It has been shown that for the same outer diameter, a straight concentric annular heat pipe outperforms a conventional heat pipe of the same outer dimensions for at least some configurations and working fluids.⁽¹³⁾

With freedom to vary the fuel element diameter and heat pipe outer radius, a wide design space exists. Unlike with the figure-eight configuration, energy deposited from adjacent fuel/heat-pipe assemblies is directly adjacent to a wick running along the heat pipe axis toward the condenser. Further, the multiple wicking pathways provide a degree of fault tolerance, particularly with regard to wick plugging and corrosion.

Examining the geometry further, however, reveals another opportunity. Note that the center of the heat pipe is empty, from the top of the fuel element to the condenser. While some of the space beyond the active fuel may be needed for axial reflector or shielding material, this space is available for other uses, including additional heat pipe segments.

Chapter 7

Multi-Stage Heat Pipe Combination

In a number of other system concepts, such as the reference reactor module for the Fission Surface Power Program,(41) adding intermediate pumped loops adds complexity, reduces reliability, increases system mass, and (because of reduced efficiency) increases thermal power necessary. Here, however, the introduction of a multi-stage two-phase heat transport system presents an opportunity to improve performance, reduce system mass, and increase reliability.

As early as 1967, the idea of subdividing a heat pipe into multiple sections had been examined,(43) in the context of reducing the effective capillary height in a heat pipe, and a multi-stage heat pipe patent was published in 1972.(31) Both inventions fall short of what is proposed here, both generally as multi-stage heat pipes, and specifically in the context of nuclear reactor heat transport. Both of these patents, while technically heat pipes using a wick for return of liquid to an evaporator, function more as a series of boilers, with heat transfer only across a plane in the axial direction.

Neither geometry serves a heat pipe reactor particularly well, as the aspect ratios are very different: the length of a heat pipe is on the order of several meters, while the diameter is typically on the order of a centimeter. The interfacial area of two heat pipes stacked as envisioned in either patent would require an extreme heat flux at the interface and would be unworkable.

Instead, here we insert an ordinary cylindrical heat pipe inside the annular heat pipe described in the previous section. The condensing region of the annular heat pipe shares a wall with the evaporator region of the conventional cylindrical heat pipe, where the




Figure 7.1: Flow diagram showing straight annular heat pipe. Diametrical wicks shown in cross-sectional view. Not to scale; ratio of heat pipe length to diameter much greater than shown.

design space permits changing the length of overlap (and with it the length of the shared condenser/evaporator space). In turn, this impacts the annular thickness of the in-core heat pipe, the resulting fuel density, power limits, and performance under failed or inoperable heat pipe conditions.

In section 7.2, a design process is presented evaluating this concept, optimizing its parameters, and comparing it to a more conventional heat pipe reactor concept.

7.1 Inverted Multi-Stage Configuration

As a brief aside, it should be noted that a multi-stage heat pipe of the sort described could also be used in a conventional heat pipe core block, where the heat pipes and fuel elements occupy alternating lattice locations. In this arrangement, a conventional heat pipe would be integrated with the core block, rejecting its heat to an annular heat pipe surrounding its condenser section. Depending on the medium to which the heat pipe assembly is rejecting heat and heat exchanger geometry, a third conventional heat pipe could be connected to reduce the cross-sectional area occupied by heat pipes. An example of such a geometry is shown in Figure 7.2.




Figure 7.2: Flow diagram showing multi-stage heat pipe in conventional heat pipe concept core block. Not to scale; ratio of heat pipe length to diameter much greater than shown.




Figure 7.3: Diagram showing dimensional parameters for multi-stage fuel-in-heat-pipe element. Not to scale; ratio of heat pipe length to diameter much greater than shown.

It is possible that the heat pipe in the core block could be reduced in size to the point it could occupy only an interstitial location, dramatically increasing the effective fuel density and reducing core size in conventional layouts. This reduction is made possible by the same impact on effective length discussed in the next section.

7.2 Design Process

Recall the relationship from Equation 2.3, describing an effective heat pipe length:

$$L_{eff} = \frac{L_{cond}}{2} + L_{adiabatic} + \frac{L_{evap}}{2} \quad (7.1)$$

Because the vapor velocity goes to zero as you approach the end of a heat pipe, the pressure drop per unit length in an evaporator or condenser section is less than that in an adiabatic section. A heat pipe with unit length where heat addition occurs over one half of that length, and heat removal over the other half of that length, has an effective length of half of the physical length of the heat pipe. By combining multiple heat pipes constantly gaining or rejecting heat energy, the effective length is reduced, and the heat transport capacity increased, subject to temperature drops across the boundaries between heat pipes. Beyond packaging considerations in the design of a system, this reduction in effective length is central to the performance gain expected in a multi-stage configuration.

For the condition where all of the heat produced in the heat pipe assembly is rejected through the condenser section of the second stage, the mass flow rate in both heat pipe stages is identical, and is calculated directly given the heat input and the heat of vaporization for the working fluid, where Q is the power and H_v is the latent heat of vaporization of the fluid:

$$\dot{m}_1 = \dot{m}_2 = \frac{Q}{H_v} \quad (7.2)$$

The total heat pipe length consists of the heated length in-core, some length between the core and heat rejection region with negligible heat gain or loss, and the heat rejection region. Assuming this length is driven primarily by system layout constraints, we can fix this length, l_{total} , ahead of exploring the parameter space. It consists of the sum of the two heat pipe lengths, less the length where they overlap:

$$l_{total} = l_f + l_{ab1} + l_{c1} + l_{ab2} + l_{c2} \quad (7.3)$$

Effective lengths, and wick and vapor areas for each of the heat pipe segments can also be defined, with subscript 1 referencing the annular heat pipe in contact with the fuel, and subscript 2 referencing the secondary heat pipe:

$$A_{w,1} = \pi \cdot [(-r_{2o}^2 + (r_{2o} + r_{wick})^2) + (r_{1i}^2 - (r_{1i} - r_{wick})^2)] \quad (7.4)$$

$$A_{v,1} = \pi \cdot (r_{1i} - r_{wick})^2 - \pi \cdot (r_{2o} + r_{wick})^2 \quad (7.5)$$

$$l_{eff,1} = \frac{1}{2} \cdot (l_f + l_{c1}) + l_{ab1} \quad (7.6)$$

$$l_{total,1} = l_f + l_{c1} + l_{ab1} \quad (7.7)$$

$$A_{w,2} = \pi \cdot (r_{2i}^2 - (r_{2i} - r_{wick})^2) \quad (7.8)$$

$$A_{v,2} = \pi \cdot (r_{2i} - r_{wick})^2 \quad (7.9)$$

$$l_{eff,2} = \frac{1}{2} \cdot (l_{c1} + l_{c2}) + l_{ab2} \quad (7.10)$$

$$l_{total,2} = l_{c1} + l_{c2} + l_{ab2} = l_{total} - l_f - l_{ab1} \quad (7.11)$$

Since the heat transfer across the wall between the two heat pipes is bounded by liquids condensing or evaporating on the two surfaces, the temperature drop across the wall can be modeled in a single dimension, dependent only on area, heat flux and thermal conductivity of the wall.

$$T_{wall,c1} - T_{wall,e2} = \frac{Q}{2\pi l_{c1} \cdot k_{wall}} \ln \frac{r_{2o}}{r_{2i}} \quad (7.12)$$

For a given set of wall properties (thickness, radius, thermal conductivity), it is obvious that as the length of overlap increases, the temperature difference decreases. If the object were only to minimize this temperature drop, the annular heat pipe would extend to the condenser of the inner heat pipe. This results in a maximum effective length for the annular heat pipe, however, which will drive an increase in the annular thickness and a reduction of effective fuel density in the core for a given power level. It is reasonable to impose minimum and maximum lengths on this section, and to optimize within that space for the best overall performance.

A starting point for determining the division of the two heat pipe stage lengths is to require that they have the same maximum power, as limited by capillary pressure. Recalling Equation 2.8, and substituting the appropriate dimensions from Figure 7.3 for each heat pipe gives the following:

$$\dot{Q}_c = \frac{\sigma_l \cdot \rho_l \cdot H_v}{\mu_l} \cdot \frac{K \cdot A_{w1}}{l_{eff,1}} \cdot \left(\frac{2}{r_{eff}} - \frac{\rho_l \cdot g \cdot l_{total,1} \cdot \cos\Phi}{\sigma_l} \right) \quad (7.13)$$

$$\dot{Q}_c = \frac{\sigma_l \cdot \rho_l \cdot H_v}{\mu_l} \cdot \frac{K \cdot A_{w2}}{l_{eff,2}} \cdot \left(\frac{2}{r_{eff}} - \frac{\rho_l \cdot g \cdot l_{total,2} \cdot \cos\Phi}{\sigma_l} \right) \quad (7.14)$$

Since the temperature difference between the two heat pipes is small, and since the wicks are made of the same porous medium they have the same permeability, these leading terms are the same on both sides of the Equation and can be excluded, leaving:

$$\frac{A_{w1}}{l_{eff,1}} \cdot \left(\frac{2}{r_{eff}} - \frac{\rho_l \cdot g \cdot l_{total,1} \cdot \cos\Phi}{\sigma_l} \right) = \frac{A_{w2}}{l_{eff,2}} \cdot \left(\frac{2}{r_{eff}} - \frac{\rho_l \cdot g \cdot l_{total,2} \cdot \cos\Phi}{\sigma_l} \right) \quad (7.15)$$

For a horizontal structure, or one operated outside of a gravitational field, the $\cos \Phi$ term is zero, and the effective pore radius, r_{eff} is the same for both wicks, leaving just:

$$\frac{A_{w1}}{l_{eff,1}} = \frac{A_{w2}}{l_{eff,2}} \quad (7.16)$$

Rearranging this result gives a relationship between the two effective lengths that depends only on the radial dimensions of the assembly:

$$\frac{l_{eff,1}}{l_{eff,2}} = \frac{A_{w1}}{A_{w2}} = \frac{[(r_{2o}^2 - (r_{2o} + r_{wick})^2) + (r_{1i}^2 - (r_{1i} - r_{wick})^2)]}{(r_{2i}^2 - (r_{2i} - r_{wick})^2)} \quad (7.17)$$

$$\frac{\frac{1}{2} \cdot (l_f + l_{c1}) + l_{ab1}}{\frac{1}{2} \cdot (l_{c1} + l_{c2}) + l_{ab2}} = \frac{[(r_{2o}^2 - (r_{2o} + r_{wick})^2) + (r_{1i}^2 - (r_{1i} - r_{wick})^2)]}{(r_{2i}^2 - (r_{2i} - r_{wick})^2)} \quad (7.18)$$

In practice, both the evaporator region of the in-core heat pipe segment and the condenser region of the secondary heat pipe segment will be constrained by heat transfer at those boundaries.

7.2.1 Comparing Multi-Stage Heat Pipe Performance to Conventional Heat Pipe

The most direct comparison between a conventional large heat pipe reactor design and the multi-stage annular heat pipe configuration is the 3-stage arrangement shown in Figure 7.2. In this comparison, the in-core evaporator of the first heat pipe segment and the condenser of the third heat pipe segment are identical to those in the conventional configuration, and we solve for the internal dimensions (evaporator, condenser, and adiabatic section lengths) that result in the highest maximum power for the assembly.




Figure 7.4: Schematic for comparison of conventional and multi-stage heat pipe.

For a fair comparison, both heat pipe designs were constrained such that they maintained the same overall length, and the same in-core evaporator and final condenser length. Wick thicknesses, permeability, and pore size were kept constant. The first and final stages of the multi-stage heat pipe have the same radii as the only stage of the conventional heat pipe. Only the lengths of the overlapping sections and the 3 adiabatic sections vary in the solution process.

Examining only the capillary limit, the solution drives toward maximizing the length of the intermediate heat pipe, and minimizing the overlapping sections. This is the result (under typical sodium conditions) across a wide range of parameters. The annular heat pipe's performance, with significantly more cross-sectional wick area, far out-performs either of the conventional heat pipe segments within it.

Constraining the comparison to include the effects of the boiling (heat flux) limit requires a minimum evaporator area on each heat pipe segment, and conduction through the heat pipe wall results in a reduction in temperature with each transition.

The maximum power as limited by boiling in the wick can be expressed in terms of heat flux, and is given by Equation 7.19. Here, r_n refers to the nucleation site radius, and k_{eff} refers to the effective thermal conductivity of a saturated wick.

$$Q_b = \frac{4\pi L_{eff} k_{eff} T_{sat} \sigma_l}{H_{fg} \rho_v \ln r_i / r_v} \left(\frac{1}{r_n} - \frac{1}{r_{eff}} \right) \quad (7.19)$$

For the purpose of the evaluation here, we can estimate the increased power required from the reduction in working temperature by imposing a fraction-of-Carnot multiplier on the heat pipe power limits.¹

Another approach to constraining the lengths of the overlapping region is to require that each evaporator section be at least as large as the in-core evaporator section. This effectively guarantees that the in-core evaporator will be limiting with respect to the boiling limit, given the non-unity power peaking factor in-core. The secondary and tertiary heat pipe segments will have a uniform temperature and azimuthally uniform heat fluxes in their evaporators.

¹In other words, when comparing the power limits one heat pipe configuration to another, the final condenser temperature is multiplied by the theoretical (Carnot) cycle efficiency for a given heat rejection temperature.




Figure 7.5: Ratio of multi-stage heat pipe limiting power to equivalent single-stage heat pipe for a specific point design as a function of length, maintaining in-core and condenser parameters and cross-sectional geometries.

Independent of the design constraint on evaporator length, the relative performance of a multi-segment heat pipe to a single, conventional heat pipe exhibits the same trend as overall length increases (see Figure 7.5). The ratio increases rapidly when overall length is first increased, but hits a sharp knee and effectively becomes constant for extremely long heat pipes (8-12 for many of the sodium heat pipe parameters examined).

For shorter overall lengths, this follows from the power optimization solution driving to the smallest allowable conventional heat pipes at either end, with the conventional heat pipe segments having a lower maximum power (limiting the total power). Since the conventional heat pipes are limiting, changing the overall length initially does not reduce its overall power limit. Once the annular segment is limiting (i.e. its length long enough that its capillary




Figure 7.6: Limiting power for multi-stage and equivalent conventional heat pipes for a specific point design. Note that while the power ratio becomes constant, total power carried by either heat pipe type declines as length increases.




Figure 7.7: Fraction of unconstrained heat pipe length assigned to annular heat pipe segment. Plot shown for sodium heat pipe with sintered wick, 1 cm diameter at 900K, with fixed 0.5m condenser (heat rejection to power conversion) and evaporator (in-core).

limit falls to the limit of the other two segments), the end sections begin to grow in length, with all 3 segments having the same power limit, decreasing as length is further increased.

The limiting power for the baseline heat pipe (i.e. conventional cylindrical heat pipe) falls immediately with increases in length, continuing to do at a steady rate. This is not surprising, nor remarkable, and is easily evident from Equation 2.8, with effective length in the denominator.

This comparison shows that for a common heat pipe reactor configuration, a multi-stage heat pipe significantly increases the amount of power that can be transported at each heat pipe location. In practice, this translates to either additional margin to failure, increased power, or a combination of both.

With the simple substitution of a multi-stage heat pipe, fabricated with essentially the same techniques and identical materials, we arrive at a *several-fold increase in the limiting power* at operating temperature.




Figure 7.8: Limiting power as a function of temperature for a conventional heat pipe (red), and its multi-stage equivalent (blue). The annular (dashed grey) and in-core (yellow) heat pipe segments are also shown, but total power is limited by the condenser section in this case.

We can also examine the heat pipe’s performance beyond a single state point, and generate performance curves showing the limiting power (and the phenomenon driving it) over a wide range of temperatures. One such set of curves, comparing the multi-stage heat pipe to a conventional heat pipe, is shown in Figure 7.8. A second comparison, with a thicker wick structure for the single-stage heat pipe and the core section of the multi-stage heat pipe, is shown in Figure 7.9. In addition to changing the wick structure, the length of each segment is again adjusted to maximize performance at the design temperature. Here, the single-stage heat pipe has a higher power limit at operating temperature, but a much lower limit as temperature declines. It does outperform the multi-stage heat pipe over a very short range of temperature,




Figure 7.9: Limiting power as a function of temperature for a conventional heat pipe (red), and its multi-stage equivalent (blue), for a similar heat pipe with a thinner wick (i.e. less capillary pumping, greater flow area). In this case, the in-core segment is limiting for the multi-stage heat pipe.




Figure 7.10: Unit cells for conventional (left) and annular (right) heat pipe configurations.

7.3 Design Features and Trades

Returning to the two-stage fuel-in-heat-pipe configuration introduced in section 7, we explore certain design features possible with the combination and their simulated performance and practical implications.

At this stage, there are three key questions to answer: is it in fact possible to increase the effective fuel density, is it possible to operate at a higher power density, and is it similarly possible to maintain or reduce the effective void fraction in-core (thus making water immersion accidents easier to mitigate)? If the answer to each question is yes, any of the existing reactor concepts can be modified for higher power, smaller size, or a combination of both.

The method that follows imposes a series of constraints such that key performance parameters remain the same. First, we require that the minimum web thickness is constant

in the two cases, assuming the limiting structural conditions are comparable. This allows a definition of pitch, P , and a calculation of total fuel and heat pipe area.²

For the conventional heat pipe configuration, it is apparent in the above relationship that maximizing heat pipe surface area, and wick area, and fuel loading occur when the outer radius of the heat pipe and fuel element are the same. This also results in the shortest conduction path from fuel to heat pipe, and a minimum web thickness allowing conduction between fuel elements (as in a failed heat pipe condition).

For each configuration, we can calculate a fuel fraction, vapor area fraction, and wick length (or area) for a unit cell. If the pitch and minimum block web thickness are the same in both cases, the result is a unit cell comparison where an annular fuel-in-heat-pipe assembly directly replaces both the fuel elements and the heat pipe lattice locations in a conventional heat pipe reactor layout.

For the conventional arrangement, the fuel radius (r_{fuel}) is simply half the pitch minus half the web thickness. For the annular configuration, the heat pipe parameters determine the remaining fuel space for a given pitch, such that the fuel radius is given by Equation 7.20, with P representing the center-to-center element spacing, and t_{web} the minimum core block web thickness between elements.

$$r_{fuel} = (P - t_{web})/2 - 2 \cdot t_{wick} + t_{vapor} + t_{wall} \quad (7.20)$$

This requires some heat pipe design, not just fuel size, to get an idea of effective fuel density, with 3 additional degrees of freedom. Recalling the 2 heat pipe limits of interest, boiling and capillary, we can impose two constraints on the annular configuration: heat flux into the annular heat pipe is no greater than the conventional arrangement, and wick area is no less.

Looking at heat fluxes in Equations 7.21 and 7.22, we see the effect of having a 2:1 ratio of fuel-to-heat pipe azimuthally in a unit cell. In the case where fuel and heat pipe radii are the same in the conventional layout, an equivalent fuel radius must transport twice the heat

²Note that in the annular heat pipe configuration, the number of fuel elements and heat pipes is equal to the number of lattice locations, while in the conventional arrangement, 1/3 of lattice locations would be occupied by heat pipes, and 2/3 by fuel elements for a 2:1 layout. Other arrangements have been studied, including 3:1 and 6:1.

Table 7.1: Minimum dimensions used for unit cell comparison of annular in-core heat pipe.

Parameter	Minimum
Web Thickness	2 mm
Clad Thickness	1 mm
Wick Thickness	0.5 mm
Annular Gap Thickness	2 mm

flux as the annular design. For the same inner wick thickness, the annular configuration has twice the margin to the boiling limit.

$$q''_{annular} = \frac{q''' \cdot r_{fuel} + t_{wall}}{2} \quad (7.21)$$

$$q''_{conv} = \frac{q''' \cdot r_{fuel}^2}{r_{HP}} \quad (7.22)$$

Ordinarily, capillary pumping power is increased by increasing the wick thickness for a given wick design. Considering that the annular design starts off with twice the wick thickness and half the heat flux, it would seem to have a significant advantage over the conventional configuration. The remaining question is whether, under reasonable dimensional constraints, the annular configuration results in a higher effective fuel density.

Suppose we establish a set of minimum dimensions for the wick, wall thickness, and the minimum web thickness for the core block, as shown in Table 7.1. Given these constraints, it is possible to solve for the pitch that gives the same fuel fraction in a unit cell. For these dimensions, a 3.53 cm pitch gives an equivalent fuel loading; increasing the pitch further results in greater loading for the annular configuration.

Comparing the remaining parameters, specifically wick length (heat flux/boiling limit) and area (capillary limit), and vapor flow area, the annular heat pipe has significant advantage: wick length and area are both approximately double the conventional heat pipe, while flow area is essentially the same. Since the flow channel is not circular, for the same area it will have a higher pressure drop and lower sonic limit. Increasing the annular gap to compensate, we still arrive at a heat pipe design with significantly better performance, while loading more fuel into the same area, provided the pitch is sufficiently large.

Having answered the basic questions establishing that the annular design can meet or exceed the key performance criteria, the remainder of this chapter discusses the additional design flexibility and impacts on performance that follow.

7.3.1 Effective Operating Temperature

Perhaps the most significant performance-related gains from moving the in-core portion of a heat pipe to an annulus around the fuel is the elimination of the core block from the conduction pathway. Instead of heat generated in the fuel passing through cladding, core block, heat pipe wall, wick, and then evaporating coolant on the inner wick surface, heat generated in fuel passes through only a combined clad/wall and wick before evaporating coolant. This difference is not insignificant, accounting for several tens of degrees under typical operating conditions, and roughly double in the case of an inoperable heat pipe.

Several benefits flow from this change. First, overall system efficiency improves with a higher-temperature link to the power conversion system. That translates to a lower reactor (thermal) power needed to meet the rated electrical output. Second, while reducing the thermal power reduces the average heat flux into a heat pipe, the uniform conduction path between fuel and wick results in an azimuthal power peaking factor of one. This may appear only as additional margin to the boiling limit under some conditions, but allows for additional throughput so long as the annular heat pipe is not capillary limited.

In the multi-stage configurations, some of this gain is offset by losses between heat pipe stages. There exists a trade space where temperature drop can be reduced by reducing the margin to heat pipe limits. This occurs as the heat flux across the boundary wall is reduced by increasing the overlapping length, which increases the effective length of the heat pipe segments.

7.3.2 Freedom to Adjust Heat Transfer Capacity for a Given Fuel Element

While not exploited in the comparison here, it is possible to flatten the power profile of a heat pipe reactor using the novel annular configuration. In an ordinary heat pipe reactor

configuration, decreasing the fuel radius decreases the linear heat rate, but increases the effective distance between the fuel and heat pipe. For a given fuel element of the annular configuration, reducing its clad radius (inner heat pipe wall and inner wick radius) increases the cross-sectional wick area and reduces power into the heat pipe, without increasing the distance between the heat source (fuel) and sink (saturated wick).

7.3.3 Power Peaking

Where the flattening of pin power peaking is desired, the annular configuration presents a degree of freedom where fuel loading can be reduced, while increasing heat transport capacity. By reducing the fuel radius, clad radius, and the inner diameter of the inner wick, the heat pipe's limits are increased. In other words, changing the reference dimension, fuel radius, moves both power generated and heat removed in the right direction in response to the change.

In a conventional heat pipe reactor configuration, reducing fuel radius also increases the distance from the fuel to the heat pipe. So while power generated is reduced, heat transport capacity is also reduced. Effectively, the only route to flattening the power distribution in a conventional layout is to resort to annular fuel pellets/rods, and likely results in the removal of more fuel to address the peaking problem than would have occurred in the annular heat pipe geometry.

Chapter 8

The Demonstration Using Flattop Fission (DUFF) Experiment

Since it has already successfully accomplished its mission, the author will discuss how he developed the plan, from basic technical concept to turning on a light with nuclear energy generated by the test.

Before discussing the opportunity for a nuclear test, and the concept it supports, it is necessary to acknowledge the co-dependence that exists between them. The concept, a very simple, few-hundred watt fission power system, directly encroaches on existing capability while offering an opportunity to reduce per-mission cost and greater capability (i.e. higher power and longer life) as the technology and design approach mature in space applications.

The state-of-the-art in space power systems, prior to the DUFF experiment, is essentially the same technology first deployed in the early SNAP program: a plutonium-238 heat source coupled to a thermoelectric generator.¹ For the last several decades, space systems have generally relied on either photovoltaic solar power, or these radioisotope thermoelectric generators. Supplies of ^{238}Pu have been an issue since domestic production at the Savannah River Site stopped more than 30 years ago. A nearly decade-long effort to restart production has not resulted in a steady supply, and even once that phase of production is reached, it only marginally addresses science mission needs: a single mission like the Mars 2020 rover consumes several years' production. Despite production for more than seven years, most of

¹SNAP 3 supplied power to a DOD satellite in 1961 using this combination.(45)




Figure 8.1: ^{238}Pu allocations in 2013.(19)

its fuel was from the existing stockpile. The U.S. Government Accountability Office reports that the next planned mission will exhaust remaining supplies.(35)

8.1 Purpose and Role of Duff Experiment

The DUFF experiment, which was the first US Government-developed nuclear power system in nearly half a century, was a test designed to explore the feasibility of nuclear testing with minimal expenditure on engineering and equipment. DUFF was the first heat-pipe reactor power system, and the first reactor power system to use Stirling engines for energy conversion. Despite this, it was not, in any way, attempting to undertake research and/or development in a technical sense. DUFF has served to substantiate claims by a number of engineers that a simple reactor system can be developed, and should not cost billions of dollars, and has served to substantially raise the profile of nuclear power systems within the NASA hierarchy. One might recall one of the questions posed during the development of DUFF:

Is the infrastructure in place within the Department of Energy complex at least minimally sufficient for the development, manufacture, and testing of a space reactor?

Unfortunately, DUFF made clear that while testing is feasible, significant leaps in technology are largely precluded by the regulatory environment currently in place. Advancing to a deployed nuclear rocket is most likely to occur through an incremental development and testing process. In light of nuclear rocket development, DUFF's purpose is in addressing the first step toward a nuclear rocket test: doing a real, simple, nuclear powered experiment. Each successive experiment in turn seeks to address a manageable number of technical and/or regulatory obstacles, with each successful operation clearing hurdles along the way to a nuclear propulsion system.

8.2 Developing the Reactor Concept Associated with Nuclear Test and its Application in Space

One of the biggest historical challenges in developing a space fission power system to flight has been the presumed need for a mission to support it. It takes an ambitious, large mission to reach a point where a nuclear reactor is the only viable energy source; such a mission also ends up with a high-profile, and becomes subject to political whims and changes in administration both of NASA and the federal government generally.

At the other end of the spectrum, missions that can be readily powered using sources other than nuclear sources generally have no interest in supporting, or even including, nuclear sources. For mission planners, unnecessary risk is to be avoided, and a new nuclear system presents both a technical and political risk, noting that the President must approve the launch of a space nuclear power system.

In the last several decades, every major effort to develop a reactor for space applications has been tied to either a high-profile mission that had no non-nuclear alternative, or lacked a funded mission in which it could be the primary energy source. Even a "low-power" system intended to expand upon existing radioisotope power system capabilities carefully avoided power levels for which radioisotope power was considered reasonable or for missions for which fuel had been allocated. The list of programs that spent millions of dollars without producing a single Watt-hour of energy is long.²

²See Table [2.1](#)




Figure 8.2: Initial Simple Fission Reactor schematic/design developed prior to DUFF test, shown alongside Advanced Stirling Radioisotope Generator concept for scale.

The system concept proposed by the author, the Simple Fission Reactor, directly challenged the long-held belief that space fission power should stake its own claim in terms of power and mission lifetime. The right fission system might not be as lightweight as a radioisotope system at certain low power levels, but reductions in launch costs and increases in costs associated with the radioisotope programs (including the failure of the Advanced Stirling Radioisotope Generator program, and limitations in the supply of ^{238}Pu) make a heavier fission system more attractive than it might have been decades ago. In other words, instead of staking its claim where it had no consistent NASA customer, a fission system needed to challenge existing infrastructure, and present NASA with an alternative that could be cheaper in the long run, while enabling missions with larger power requirements.

The initial Simple Fission Reactor target was 500W of electric power, a level where its mass would be comparable to a radioisotope system, with the stated goal of being able to supply power to missions with lesser power needs that could accommodate slightly higher mass. This low power level, which is lower than what was proposed by the *Small Fission Power System Feasibility Study Final Report*([29](#)) a few years prior, enabled several important design features allowing for reduced mass and reduced system complexity.

As shown in Figure [8.3](#), the concept consists of a nearly cylindrical highly-enriched uranium core (in red), with a central control rod, reflector, and shield comprising the nuclear components, with heat pipes connecting Stirling converters through the shield to




Figure 8.3: Initial Simple Fission Reactor schematic/design developed prior to DUFF test, showing components of design.

the beryllium oxide reflector. The reduced power level relative to what was previously considered the lower bound of fission systems is what enables the heat pipes to be located in the reflector, not the fuel, which reduces the fuel mass considerably. In reducing the core size, in particular its diameter, the shield volume is significantly reduced. The reflector thickness is relatively unchanged in relocating the heat pipes, but again, the smaller core diameter reduces its mass.

Contact conductances, or the degree to which energy is transferred via conduction at a material interface (in this case core-reflector), would ordinarily present some cause for concern with this approach. However, heat fluxes are sufficiently low across the boundary that the temperature rise is small enough to preserve reasonable power conversion efficiency at temperatures easily tolerated by metallic fuel.

The design was only carried so far as to show that a lower power, lower mass, lower technical risk system could be developed to compete with certain radioisotope-powered missions, and most importantly, that such a system utilized technologies and materials that existed and might lend themselves to testing. Initial reception to the idea, even among some

of biggest proponents of space fission power systems, was cold. It had been studied before, and was not viable.

As a matter of fact, the idea of a small radioisotope-replacement nuclear fission power system had been studied, with a report issued in 2010.⁽²⁹⁾ It sought to examine the feasibility of a system that could produce 1kW of electrical power over 15 years, and be ready for a 2020 launch. While the study's "ground rules" set about minimization of system mass and volume, the team recognized the importance of practicality and cost: "At the same time, the team judged that low cost and low risk (both developmental and operational) were even more important than mass."

The study argued that a system at a lower power level, such as $500W_e$, would "result in a relatively small mass difference," because the core is "already near the minimum size needed to sustain a fission reaction." The system concept recommended by the study made use of U7Mo fuel with 12 heat pipes dispersed within it. It was surrounded by a BeO reflector, and used thermoelectric generators for power conversion. The core assembly mass was estimated at 133kg, at a thermal power of 13kW.

That study gave passing consideration of Stirling engines, but with several members of the team lifelong proponents of thermoelectrics, it should come as no surprise that they became the conversion technology of choice. As discussed later, selecting the right team is important.

8.2.1 Design and Analysis of the Simple Fission Reactor

Several principles guided development of the Simple Fission Reactor concept. The goal was to make design decisions so that a credible, attractive system did not require extensive analysis up front. Selection of newer (or perceptually less mature) technologies like the use of Stirling engines was based on engineering judgement. If there was a technical issue that could be resolved or averted by making a design change, quickly using engineering judgement, instead of slowing down to analyze or optimize, make the design change. With that approach, one reaches a design like the one proposed as the Simple Fission Reactor very quickly.

Nuclear Design

The attractiveness of the system comes from the reduced system mass as compared to a 1kW_e system as previously studied where the heat pipes are integral to the fuel. Does moving the heat pipes to the core periphery result in a reduction in fuel mass? Yes.

By moving the heat pipes to the core periphery, a nuclear model of the system becomes very similar to a simple reflected sphere. The heat pipes can only absorb neutrons that would otherwise escape the core, and the nearly empty vapor channel no longer opens a leakage path. They can be excluded from the crude model with little impact on the determination of critical size. A two-cell MCNP model provides a first-order estimate of reactor mass and size, and gives credit to the idea that a smaller fission system could compete on a mass basis with radioisotope systems. Less than 20 kg of uranium metal fuel, and 40 kg BeO reflector mass results in a critical system.

Depletion/Fuel Burnup

One will notice in examining most of the literature discussing space fission power systems that very little time is given to the topic of depletion. In some instances, this reflects concept immaturity; for most missions, power levels, and system lifetimes, burnup has a relatively small effect on the system design.

For this 500W_e system, we can quickly dismiss any need to calculate burnup or its effects over time, simply based on engineering judgement and past experience with nuclear reactor systems of this sort. Nonetheless, here we conduct the mathematical exercise to substantiate that judgement, calculating it based on an expected lifetime and approximate fissile mass.

Assume a system efficiency of 20%, and a lifetime of 10 years. That implies a thermal power of 2.5kW , and a total of $220,000\text{kWh}$ of thermal energy produced. Assuming 200MeV per fission, each fission produces 9×10^{-15} Wh. Dividing the total thermal energy produced by the energy recovered per fission gives the total number of fissions, here 2.5×10^{22} . That number is approximately 0.04 mol, or 10g fissioned over 10 years. Fuel loading in this small system will be in excess of 10kg, meaning that burnup will be less than 1/1000th of the available fissile inventory.

Thermal Modeling

Ultimately, the goal of thermal modeling here is to show that a reasonable hot end working temperature for power conversion can be reached without excessive fuel temperatures. Intuitively, at the low power level of this system, we expect the temperature rises to be small.

Rough calculations were performed, and showed a thermal system so benign that more detailed modeling was not conducted. These calculations included a 1-D spherical temperature rise model in the fuel and calculation of the temperature drop across a contact conductance in vacuum.

The difference between the fuel surface and peak fuel temperature is given by the relationship below, assuming approximately constant thermal conductivity:

$$\Delta T_{fuel} = \frac{Q}{8\pi R_{fuel}k} \quad (8.1)$$

Incorporating rough estimates of the relevant parameters gives the following:

$$\Delta T_{fuel} = \frac{2500W}{8\pi(0.2m)(25W/mK)} = 20K \quad (8.2)$$

For a system expected to operate at a peak temperature on the order of 1000K, a difference of 20K between the peak and surface fuel temperature is of no concern. Similarly, the temperature drop through a few-centimeter thick reflector made of beryllium oxide is even less concerning, as it has a thermal conductivity an order of magnitude higher than uranium, and smaller heat fluxes than in the fuel.

The interfacial conductances between fuel and reflector, and between reflector and heat pipe, can be estimated assuming a static air gap and Fourier's law for operation in a normal atmosphere. Even assuming a factor of 100 reduction for operating in vacuum of space, with no other treatment, the temperature difference is small.

Summary

It should be noted at this point that the Simple Fission Reactor concept is not some optimized design that results from a trade study. It could no doubt be modeled more accurately, optimized to reduce mass further, improve system efficiency further, or scaled up or down in power to match a particular need. Those efforts have a place, but these initial calculations are sufficient to advance the concept and testing.

8.3 Identifying the Opportunity for a Nuclear Test

A few months prior to the development of the Simple Fission Reactor concept, funding levels within NASA and DOE for space fission programs were at fairly low levels, and research was more focused on basic science than system development. While visiting the Nevada Nuclear Security Site studying the possibility of conducting cross-section measurements using the zero-power critical experiment machines recently moved there from Los Alamos, the author inquired as to what power level the zero-power facility was licensed (a critical system having some number of fissions occurring, and hence some energy being generated). Its license was not to a power level, but instead a maximum dose rate—planting the seed for what would become the DUFF experiment.

The basic idea for a test was simple: utilize an existing "zero-power" critical experiment machine with a reactor configuration not completely different from what was proposed in the concept, connected to heat transport and power conversion, in the form of heatpipe and Stirling converter. Perform the test as quickly and quietly as reasonably achievable, but perform the basic—and novel—task of producing electrical power from nuclear fission, and turn on a light. In spite of the many prior efforts operate a new space nuclear power system, the most recent to that point was SNAP-10A in 1965. The Department of Energy, since its inception, had never operated a non-defense nuclear power system.³

³Naval Reactors, while reporting to both the Department of Energy and the Department of the Navy, is responsible for naval reactor power plants. Outside of the power systems described herein and the power power systems used in the nuclear fleet of submarines and aircraft carriers, no DOE managed/operated nuclear system has produced electricity since DOE was formed.

With this approach, the goal was to minimize regulatory hurdles, dial back technical complexity, and set the bar for real nuclear systems separate from decades of simulation and speculation.

8.4 Assembling a Team: Strategic Partners, Organizations, and Individuals

At this point in the development of the first experiment, the author's motivations for taking it on bear some explanation. On the heels of participating in the response to the earthquake, tsunami, and nuclear accident at Fukushima-Daiichi, the author saw the nuclear industry as a whole at a crossroads. There was more than just a scientific or engineering question to answer, but one that goes to a deeper purpose, and one looking for a signal as to the future of that industry: was the nuclear renaissance real, or would the industry head into a long, slow decline? Assembling a team to successfully conduct the DUFF experiment meant assembling a team of people willing to ask the same questions, and take the necessary risks to answer them.

8.4.1 Necessary Team Qualifications

Obviously, simply wanting to do a ground nuclear test of a space reactor concept is not enough—there's a long list of retired engineers who spent whole careers trying to get to that point. The core team needed to be one with the right technical skillset—overlapping and complementing each other—supported by project management, political, and marketing skill to make it happen. It also needed to be small, and non-competitive. There was no team-lead, or project manager, in the conventional sense. It was by all accounts a team.

Many of the ordinary technical skills are fairly easy to identify. Nuclear design, mechanical design, power system design, and system integration are obvious. Each of the core team members could be considered among the best in at least two of those design areas. The team also needed members who could fend for themselves as far as funding—and were willing to incur sometimes significant expenses personally to keep things going.

What is often overlooked in the interest of safety, programmatic development, or career stability is the ability to exercise engineering judgement, and to do so with the rest of the team's trust. It's incredibly easy as a subject matter expert, particularly within government programs, to engage in expensive study when experience and judgement are capable of rendering a "good enough" answer to a technical question. The team must be able to collectively assess risk, and focus effort on the challenges where that effort is likely to have the greatest return. That means being able to trust that the team member that looks into an issue correctly assesses its significance or how to address it; it also means knowing the design and plan well enough to be able to quickly raise an issue or put one to rest. The team was kept purposefully small, with four members including the author.

8.4.2 Necessary Organizations

Given that this is a system test involving a critical experiment, highly enriched fuel, and is intended for a space application, three organizations would have to come together to make it work. NASA is an obvious participant for the space application relevance of the experiment; however, the involvement of NASA as an organization was much more specific. NASA-GRC had both the heat pipes and Stirling converter hardware that would eventually be used on the test. Los Alamos National Laboratory beyond including the author, had the concept, design expertise, and tools to support the experiment, as well as operational control of the critical experiment facility. Lastly, the Nevada Nuclear Security Site housed the critical experiment facility, and would play a role in making the test actually happen.

8.4.3 Strategic Exclusion of Certain Entities

A careful reader will notice that a key participant, if not gatekeeper, for nearly all prior efforts to develop a space fission power system, is absent from the described team. The author decided early on to limit the number of people aware of the effort's existence, in part because of its competition with existing funded activities, but to explicitly exclude the Department of Energy's Office of Nuclear Energy. Its habitual burdening of space reactor projects with administrative studies was capable of crippling some of the best efforts, and

in some cases managed to justify spending more money answering the question of whether to study an issue than would have been spent to actually study it as proposed.

Ordinarily, the funding of an activity like the DUFF test by an organization like NASA would have triggered DOE-NE's involvement through the Work for Others processes within DOE and the national labs. DOE Order DOE O 481.1C specifically requires that DOE/NNSA resources are made available to non-DOE/NNSA entities, which includes NASA. That same order, in section 4.o., requires that "Work that involves a space nuclear reactor or non-commercial power reactor and radioisotope power source projects at non-NNSA facilities requires the concurrence of the Director, Office of Space and Defense Power Systems. Such work, when performed at an NNSA facility, requires coordination with the NNSA Office of Institutional and Joint Programs." (9)

It should also be noted here that one of the informal (but inflexible) requirements on the way to concurrence from DOE was that the "space nuclear reactor" development work must include heat transfer and power generation. This "all-DOE" approach significantly increased the outlay to DOE, including the management fees for DOE itself and the prime contractors at the national labs. NASA was being strong-armed into paying DOE for work it could do in-house.

The Work for Others requirement immediately creates two problems for DUFF. First, it adds two additional approval steps to any externally-funded effort. By itself, that manifests mostly as a barrier in time and likely delays operation of the experiment. Second, and perhaps most significantly, in putting the Office of Space and Defense Power Systems in the approval process, it creates a conflict of interest that acts as a barrier to any nuclear reactor technology that seeks to compete with radioisotope power systems. This office within the Department of Energy is essentially funded in full by efforts to supply radioisotope power systems, managing over \$120M(33) in NASA-funded infrastructure that would be rendered unnecessary, at least in part, by a successful power system of the sort the Simple Fission Reactor and DUFF experiment would make available. Many of the failed attempts to develop reactors for space applications have failed under a strategy of protectionism, often requiring the participation of multiple labs and external hand-picked consultants not needed to complete the work proposed by a NASA sponsor. The tactic results in underwhelming

system concepts, delivered for a high price tag and without a viable path to building, testing, and flying one. Even the *Affordable* Fission Surface Power Program carried a price tag over \$1 billion dollars.

Avoiding that means avoiding any transfer of funds between organizations, requiring each of the participants—LANL, NASA-GRC, and NSTec—to fund their respective portions with internal research and development funds. Of course, those funds can be scarce, and highly desirable within each of those organizations. The case had to be made that the internal contribution would result in benefit for the individual institution, but would also mean that funding on a tight time scale and from internal sources would be quite limited. As luck would have it, all three organizations were ultimately able to make it work, but not without some significant risks on the part of the core team.

8.4.4 Getting Started, Finding Funding

Ordinarily, how one finds funding would not be discussed in this context beyond an acknowledgement. Responding to calls for proposals, an evaluation process (usually by some form of committee), and awarding of a grant is typical. Informal discussions amongst what would eventually become the DUFF team developed into a test plan with a simple goal: less than one year to a working demonstration. This plan was presented by the author in early February, 2012, to a group of interested parties at LANL.(Dixon et al.)

That presentation was well received by some, but not all in attendance. Surely, the Simple Fission Reactor concept and an associated nuclear demonstration got a positive response, but the proposed schedule, budget and funding approach were ridiculed by line management. Neither the idea of an actual nuclear demonstration, DUFF, within a year, nor performing such a demonstration for less than \$2 million (as presented at the time) was reasonable. Development of a proposal was suggested. NASA was the obvious source of funding and beneficiary of the development of the system. A committee could evaluate it, and the DOE Office of Nuclear Energy could oversee it and the funds transfer to the lab. As with past programs, if funded, a paper reactor was the likely product.

Instead, the case was made that this be taken on with internal funding, demonstrating to the team that the lab had a real interest in reactor design and development, to NASA

that the lab (with DOE) was capable of doing nuclear experiments, and that there was enough value in the concept to invest internal funds. A test of this type, funded with NASA dollars, to convince NASA that a space reactor was possible does not make sense. The DUFF experiment was not going to be a flight system that NASA could use, it was only a proof-of-concept device. DUFF is analogous to what inventor might use to convince a venture capital group to invest in developing a product, not a product intended for sale to its customer.

More specifically, the author sought to add value to the demonstration through the use of internal funding. The message to NASA would be along the lines of “this is easy enough and cheap enough that we went ahead and built one for you to see.” But it also needed to stand out from the billions previously spent on space fission power system studies and development efforts. If one compared its specifications to certain prior system designs, its performance might be underwhelming.

One of the problems in making the case that this demonstration would be significant was the perception that there has been ongoing, meaningful, decades-long progress in the space reactor and DOE nuclear communities. It came as a surprise to many in management that the last space reactor system test was half a century ago. Despite the U.S. Government undertaking construction of over 100 research reactors (i.e. not including those used for naval power and propulsion), none had been undertaken since the formation of the Department of Energy. And among all existing reactors at DOE facilities, none had been used to produce electricity. It was necessary to bring the audience, including NASA as a potential customer, back to reality, bleak as it may be, before it could be recognized that the production of electricity, with energy from nuclear fission, would be fundamentally new.

This problem is far from unique to space reactors within federal research and development organizations. With an eye towards the publicity associated with the next great thing, these groups all too often try building the next level before the current one sets up. We need a vision of the tall building, but cannot leave out critical engineering elements simply for trying to be first at the top.

With interest (but not funding) at the two institutions, LANL and NASA took the case to Nevada in early March, with representatives from LANL, NASA-GRC, NASA-MSFC, and

the Jet Propulsion Laboratory in attendance. NSTec, operator of the site, was very receptive to the idea of testing at their site, and committed to covering their staff costs in support of the test. In fact, each institution covering their organization's costs proved to be a key component of the success of the project. In addition to the complications of exchanging funds between organizations, it would have involved many people not essential to the execution of the test. This also skipped the usual steps developing task plans, and piling up paper on organizational roles and responsibilities. As far as this team was concerned, there were no titles or individual tasks, it was more of a collective workload with obvious leads for certain pieces, with the rest addressed by whomever was most available.

Getting funding for the experiment ultimately required a stroke of luck, or miscommunication, depending on one's point of view. The author traveled to present a couple of conference papers but with the primary goal of facilitating a call between a NASA program director with the authority to fund NASA's part of the project, and an associate director at Los Alamos National Laboratory. As it unfolded, LANL ultimately agreed to fund the experiment because NASA was committed, and vice versa. The enthusiasm on NASA's part with regard to skipping the usual DOE-NE/NASA structure for the project was impossible to miss. The key commitments were made, but it would be nearly two more months before funding was fully in place and the process of designing the experiment could begin.

Along the way, there were a number of efforts to dilute what, so far, had been a very simplistic definition of *success*. The author saw failure to deliver power to a load as failure of the project. Surely, one can learn from failure, but failure is not the objective. *Success* had become synonymous with *good effort*, not *actual* success. The team made clear, over numerous objections, that success meant illuminating a light with power generated by the system, nothing more or less. The hope was to collect some useful data proving modeling capabilities and nuclear data, but efforts to do so *could not* slow the process.

8.5 Advancing the Design, and Preparing for the Test

The first question to resolve in advancing a design was determining if any existing critical experiments could be used as the basis for the test. At the time, the National Criticality

Experiments Research Center (NCERC) had recently finished a move from Los Alamos, New Mexico, to the former Nevada Test Site, now known as the Nevada National Security Site, north of Las Vegas. That process included modernization of some control systems, new safety analysis, and upgraded but not well-exercised hardware using existing nuclear material and mostly existing mechanical components.

8.5.1 Identifying Candidate Experiments

Initially, the plan included manufacturing fuel and reflector material (tasks that would later become part of the KRUSTY demonstration), but that plan would evolve during the month of February to make use of existing stock: first with consideration of the Rocky Flats shells, then ultimately Flattop. Flattop was an ideal candidate: a small, highly-enriched uranium core very similar in size and dimension to the concept. It was nearing operational status, and had been previously run up to pedestal temperatures of 300C. It's a highly-reflected, fast-spectrum system. And it just happened to have a hole through to the center of the core just over half an inch⁴ in diameter.

In late March, with Flattop chosen as the critical experiment machine to be used in the demonstration, Dr. Poston coined a name for it: Demonstration Using Flattop Fission, or DUFF. In his words, "DUFF continues with the long tradition of Simpsons' names, and is also appropriate because David Duff Dixon gave us the kick in the pants to start thinking about doing something at DAF near-term."

8.5.2 Heat Pipe Selection/Design

Glenn Research Center had been performing life testing on a set of heat pipes, shown in Figure 8.5, that were exactly half an inch in diameter, and had been previously tested up to 300C. One of these heat pipes was removed from the life test and sent to Nevada (though it would not be the heat pipe actually used for the test).

Sunpower helped to identify and deliver an existing set of Stirling converters suitable for low temperature operation (Figure 8.6). The Buzz converters, which had not been run

⁴0.505" per drawings, but this dimension had not been reliably verified in some time.




Figure 8.4: Flattop Schematic Drawing.



Figure 8.5: Heat pipe testing at Glenn Research Center, 2008.



Figure 8.6: First operation of Buzz converters on test stand, 90 days prior to experiment




Figure 8.7: Buzz Converters with Copper Heat Pipe Mounting Block

in years, took some effort to get going. GRC then designed and fabricated a copper block for mating the engine pair to the heat pipe (Figure 8.7). While GRC was readying the power conversion hardware, LANL was working through the nuclear design and approval process for getting the hardware into the facility and for configuring Flattop neutronically, developing a system model for the experiment. LANL staff at the test site were also working to address a control system issue that had Flattop in a non-operational state. Almost all of the preparation work took place over little more than two and a half months.

8.5.3 Facilities and Experiment Preparation

Introduction of a power conversion system that would take heat from Flattop operable at the relatively low temperatures that would be available during a Flattop run requires a low heat rejection temperature to the environment. Getting to these low temperatures ($< 0^{\circ}\text{C}$) in space is trivial, but in a room temperature environment requires the use of a chiller.

This led to a need for power for the chiller, something ordinarily trivial in any normal experimental or industrial environment: either 20A 240VAC circuit for a chiller that would be able to maintain the targeted cold-end temperatures, or a 20A 120VAC circuit for a smaller model (that would not maintain as low of a cold end temperature). A series of simulations of the experiment performed by another team member showed that the smaller chiller might be able to hold a low enough temperature to produce electricity with the Stirling engines, but that the result was heavily dependent on the rate of reactivity insertion and a number of other assumptions. With the larger chiller, there was a much larger parameter space that would still allow the Stirling engines to operate. Existing 120VAC circuits supplying receptacles were available, but a receptacle to plug in the larger chiller was not.

Less than one month prior to operation of the experiment, management was pressuring the team not to “rock the boat” by pressing for installation of a 240VAC receptacle. Test site personnel presented the task as an extraordinary undertaking, something that would take many months to accomplish. Somehow, it was not until this news reached the author that the trivial nature of the task was outlined: service panels were adjacent to the chiller location, with available capacity, needing less than an hour of effort and less than \$100 in parts from any home improvement store or electrical supply house. It was a task that the average electrician performs as a matter of routine, just not in this particular government facility.

The trivial nature of the actual work to be performed, combined with the risk to the experiment’s success using the smaller chiller, proved too much for the bureaucracy to justify. With some effort, a path in existing policy was identified to install the receptacle as a temporary facility modification, remaining in place until it could at some future (far distant) date be included in facility drawings with all of the usual bureaucratic processes completed. Immediately following DUFF, this same process was used to advance several other minor electrical tasks, one of which had been in process for over a year (now completed within a couple of weeks).

Data Connection for Power Conversion System

The wireways connecting the experiment location to the control room are not readily accessible for running key cabling between the two locations. Several coaxial cables were available, with one set connected to a pair of consumer-grade Ethernet-over-Coax adapters for remote connection to the computer adjacent to the experiment monitoring the Stirling converters. A small control box for simple control of the Stirling converters was also connected via these cables.

Visual Indicator

One of the author's early objectives was to generate electricity from fission and turn on a light. Given the expected power level, GRC prepared an electroluminescent (EL) panel, which would sit next to the converters. When switched on (with one of the aforementioned cables dedicated to this function), it would be visible on a closed-circuit video feed in the control room.

At some level, this light was a gimmick, in a way an attempt to recreate the iconic photo of the string of light bulbs powered by EBR-I. A digital readout with a dummy load would have been perfectly adequate for testing the power conversion system and the experimental setup as a whole. It would not have resulted in the same image, or been as impactful as a blinking light.

As it turned out, during execution of the test, it performed a critical function. Connectivity to the instrumentation rack was lost during startup, attributed to the radiation environment affecting the Ethernet-over-Coax adapters. The mechanical switch and light panel provided the only live indication that the Stirling engines were operating, producing electricity, and delivering it to a load.

8.5.4 Test Assembly and Operation

Most of the core team (6 people) of “visitors” arrived at the Device Assembly Facility on Tuesday, September 11, to begin uncrating and assembling all of the test hardware. This included connecting to facility instrument wiring and power, installing computer equipment,




Figure 8.8: Flattop with Heat Pipe Inserted (end of heat pipe visible as small silver circle in center of assembly)

helium fill gas and a chiller for the converters, moving and installing the proper nuclear material (Flattop has three different cores and numerous fuel mass adjustment pieces), and numerous other tasks. The heat-pipe-and-Stirling-converter setup was tested with an electrical heater. Wednesday brought installation of the heat-pipe-and-Stirling-converter setup into Flattop, and an approach to critical.

Operation of the test on Thursday, September 13 accomplished the basic goal of illuminating a light panel with energy from Flattop, but was a less-than-perfect test. Invited guests would be present on the following Tuesday, so the team had to act quickly to make modifications to ensure the demonstration worked as planned.

As mentioned earlier, the connection to the system collecting Stirling engine data was lost during the first run. That failure coincided with increased Flattop power, and the presumption at the time was that no data had been collected. Fortunately, some hours later a connection was re-established, and the complete set of Stirling engine data, including all of the measured heat pipe data, was recovered. Since the computer had not failed, focus shifted



Figure 8.9: DUFF team in control room on September 13, after first DUFF run.

to the Ethernet-over-Coax adapter and the cabling to the control room. The Ethernet-over-Coax adapter endpoint was replaced and relocated to reduce its radiation exposure. On the second run, which sought to produce power over a longer duration (in time), reactivity was inserted at a slower rate. Figures [8.11](#) and [8.12](#) show the temperatures, power, and electrical output from the two runs.



Figure 8.10: DUFF setup and attendees for second DUFF run.




Figure 8.11: Reactor power, temperatures, and Stirling engine output from DUFF first run.




Figure 8.12: Reactor power, temperatures, and Stirling engine output from DUFF second run.

Chapter 9

Conclusions

One of the challenges faced by many large research and development programs is that they often do not survive to completion. This can be for any number of reasons, including technical challenges, cost overruns, delays, scope or mission creep, but it is also often for either lack of publicly visible progress or changes in high level priorities. Many prior attempts to deploy a space reactor system, including SP-100, TOPAZ (U.S.), JIMO, FSP, and many smaller programs, have been cancelled as the true costs of development come into focus.

In each of these programs, cancellation occurred after the spending of over \$500M or development estimates growing over \$1B. Unfortunately, while significant design and development resources were expended, none of these programs truly advanced the state of the art. Concepts were envisioned, designs proposed, and specifications written, but none ever put a neutron to work producing electrical power or thrust. What these programs have done instead is to create a picture of hypothetical performance that makes the sale of more reasonable system that might actually make it to an operational test much more difficult. The incremental approach presented here makes sure to both advance the state of the art and leave behind useful facilities and/or test articles to benefit both the program at hand and nuclear research more generally. This has to happen at each phase, such that if the following phase never happens, it is still a complete and useful task.

Key in successfully completing the test in this time frame were four factors: (1) organizations sufficiently committed to the test objectives so as to fund their own efforts internally, (2) an aggressive, if not almost impossible (by some) timeline, (3) a simple, clear,

and nonnegotiable test objective (the illumination of a light with energy from fission), and (4) a capable team focused on that objective, not on individual tasking. Each organization coming up with internal funding for its part of the experiment was a key element to success. Each institution had a vested interest in the experiment’s success, as well as the scrutiny should it not have worked.

The aggressive timeline was also important from the standpoint that it precluded business-as-usual trade-study efforts, and mandated prompt decision-making. It forced the participants to constantly ask whether a particular issue was of a high enough priority to be deserving of further effort. As a result, effort was generally well-focused on the highest priority tasks (the “long poles” in the tent) in a manner that did not overlook other “details” that were lower priority but nonetheless important.

The nature of the objective, the illumination of a light with energy from fission, is not the outcome of some exercise in defining “success criteria.” It’s simple if the light lights, it is a success, if not, failure. That is not to say that a success means that everything was perfect, which is rarely the case, or that failure means nothing was learned (the most tragic of failures). It was a goal that everyone involved clearly understood, and that helped to focus efforts on even the most trivial of activities. When making decisions, at the forefront was whether that decision would help or harm the odds of success.

An effective team was obviously key here, both amongst the group that was at the Device Assembly Facility getting hardware in place as well as with the extended team. It bears mention that just in the photo from test day, Figure 8.10, 10 institutions are represented. Without interest, input, and support from all of these people, the test wouldn’t have been possible. Even then, there were setbacks, and a lot had to fall in place just right to finish on time.

More broadly, the advances in heat pipe design shown by the model presented opens doors to higher power heat pipe reactor systems. Slight increases in heat pipe complexity significantly simplify overall system complexity, while manufacturing at the system level is simplified through the use of multi-stage heat pipe designs. These designs represent a novel embodiment of a heat pipe, but one that is nonetheless easily reduced to practice.

Finally, the effort as a whole demonstrates an ongoing need to continue doing real new things to keep skill set alive (with respect to reactor engineering, testing, and licensing), and to continue drawing smart capable engineers into the nuclear field, in order to ensure that the body of knowledge continues to grow and advance the state-of-the-art.

Bibliography

- [1] Amiri, B. W., Kapernick, R. J., Sims, B. T., and Simpson, S. P. (2007). A new capability for Nuclear Thermal Propulsion design. In *AIP Conference Proceedings*, volume 880, pages 438–448. [116](#), [117](#)
- [2] Boo, J. H., Park, S. Y., and Kim, D. H. (2005). An experimental study on the thermal performance of a concentric annular heat pipe. *Journal of Mechanical Science and Technology*, 19(4):1036–1043. [59](#)
- [3] Burkes, D. E. (2007). An Overview of Current and Past W-UO₂ CERMET Fuel Fabrication Technology, INL/CON-07-12232. In *Space Nuclear Conference*. [15](#)
- [4] Chi, S. W. and W., S. (1976). Heat pipe theory and practice. *hemi*. [21](#)
- [5] Commission, A. E. (1955). *Reactor Handbook: Engineering*. McGraw-Hill. [4](#)
- [6] Commission, U. A. E. (1970). Nuclear Reactors Built, Being Built, or Planned in the United States As of December 31, 1975. Technical report, Division of Technical Information Extension. [6](#)
- [7] Congress, U. S. (1946). Atomic Energy Act 1946. *Public Law*, pages Pub. L. 79–585. [5](#)
- [8] Cotter, T. (1965). Theory of Heat Pipes. Technical report, Los Alamos Scientific Laboratory of the University of California. [21](#)
- [9] Department of Energy (2011). WORK FOR OTHERS (NON-DEPARTMENT OF ENERGY FUNDED WORK). [91](#)
- [10] Dixon, D. D. (2006). Design of a 25-kWe Surface Reactor System Based on SNAP Reactor Technologies. In *AIP Conference Proceedings*, volume 813, pages 932–941. AIP. [12](#)
- [11] Dixon, D. D., McClure, P., Beller, T., Poston, D., Gibson, M., and Bagg, S. (2013). The DUFF Experiment: Why and How. In *Nuclear and Emerging Technologies for Space, American Nuclear Society Topical Meeting*. [92](#)
- [12] EDF Energy (2018). Taishan 1, world’s first EPR connected to the grid — EDF Energy. [15](#)

- [13] Faghri, A. and Thomas, S. (1989). Performance characteristics of a concentric annular heat pipe: Part I-Experimental prediction and analysis of the capillary limit. *Journal of Heat Transfer*, 111(4):844–850. [60](#)
- [14] Finseth, J. (1991). Overview of Rover Engine Tests, NASA-CR-184270, February 1991. Technical report, Sverdrup Technology, Inc. [15](#)
- [15] Gaugler, R. S. (1942a). Capillary heat transfer device for refrigerating apparatus, 2,448,261. [17](#)
- [16] Gaugler, R. S. (1942b). Heat transfer device, 2,350,348. [4](#), [17](#)
- [17] Gebicke, M. (1992). The SP-100 Nuclear Reactor Program, Should It Be Continued? [10](#), [11](#)
- [18] Gerrish, Harold P., J. (2014). Nuclear Thermal Propulsion Ground Test History: The Rover/NERVA Program. [14](#)
- [19] Green, J. L. (2013). NASA ’ s Planetary Science Program Overview. [xi](#), [81](#)
- [20] Grover, G. M. (1964). Method and apparatus for heat transfer, 3,229,757. [18](#)
- [21] Haslett, R. (1995). Space Nuclear Thermal Propulsion Program Final Report. Technical report, Grumman Aerospace Corporation, Bethpage, NY. [14](#)
- [22] Kapernick, R. J. and Dixon, D. D. (2008). Comparison of methods for evaluating nuclear thermal propulsion tie-tube designs. In *AIP Conference Proceedings*, volume 969, pages 430–438. AIP. [128](#)
- [23] Kessler, W. E., Cordes, O. L., Dinneen, G. A., Fletcher, R. D., and Brockett, F. (1965). SNAPTRAN 2 / 10A -3 DESTRUCTIVE TEST RESULTS. Technical report, Philips Petroleum Company. [9](#)
- [24] Kim, K., H.B., K., and Lee, K. (2003). Development of Heat Pipe Heating Roller in Laser Printer. In *Proceedings of The 7th International Heat Pipe Symposium, Jeju, Korea, pp. 65– 69*. [59](#)

- [25] Kingsbury, N. R. (1992). Space Nuclear Propulsion: History , Cost, and Status of Programs. [15](#)
- [26] Kirkpatrick, M. E. (1972). ANNULAR HEAT PIPE. [59](#)
- [27] Koenig, D. R. (1986). Experience Gained from the Space Nuclear Rocket Program (Rover). Technical report, Los Alamos National Laboratory. [28](#)
- [28] Lichtenberger, H. V., Novick, M., Cerutti, B. C., Cameron, R. A., McGinnis, D. F., Pettitt, E. N., Whitham, G. K., Haroldsen, R. A., and Kock, L. J. (1953). Progress Report on the Experimental Breeder Reactor. Technical report, Argonne National Laboratory, Lemont, IL. [6](#)
- [29] Mason, L., Bailey, S., Bechtel, R., Elliott, J., Fleurial, J.-P., Houts, M., Kapernick, R., Lipinski, R., MacPherson, D., Moreno, T., Nesmith, B., Poston, D., Qualls, L., Radcl, R., Weitzberg, A., and Werner, J. (2010). Small Fission Power System Feasibility Study Final Report. Technical report, NASA. [83](#), [85](#)
- [30] McClure, P. R., Dixon, D. D., Poston, D. I., Kapernick, R. J., Reid, R. S., and Dasari, V. R. (2014). Mobile heat pipe cooled fast reactor system. [ix](#), [42](#), [46](#), [47](#)
- [31] Moore, R. D. (1970). United States Patent Moore, Jr. 54 SEGMENTED HEAT PHPE. Technical report. [61](#)
- [32] NASA (2019). Public Release Summary of the Geostationary Operational Environmental Satellite-17 (GOES-17) Advanced Baseline Imager (ABI) Loop Heat Pipe (LHP) Mishap Investigation Board Report. Technical Report July, National Aeronautics and Space Administration. [22](#)
- [33] NASA (2020). FY 2021 President’s Budget Request Summary. [91](#)
- [34] Newhouse, A. (2003). Forum on Concepts and Approaches for Jupiter Icy Moons Orbiter Science Capabilities & Workshop Goals. [11](#)
- [35] Oakley, S. S. (2017). Improved Planning and Communication Needed for Plutonium-238 and Radioisotope Power Systems Production GAO-18-161T. [81](#)

- [36] Poston, D. and Kapernick, R. (2012). Design and analysis of megawatt-class heat-pipe reactor concepts. In *International Congress on Advances in Nuclear Power Plants*. [52](#)
- [37] Poston, D. I. (2001). The Heatpipe-Operated Mars Exploration Reactor (HOMER). In *AIP Conference Proceedings*, pages 797–804. [46](#)
- [38] Poston, D. I. (2018). Nuclear testing and safety comparison of nuclear thermal rocket concepts. *Nuclear and Emerging Technologies for Space, NETS 2018*, pages 345–348. [39](#)
- [39] Poston, D. I., Gibson, M. A., Godfroy, T., and McClure, P. R. (2020). KRUSTY Reactor Design. *Nuclear Technology*, 206(sup1):13–30. [7](#)
- [40] Poston, D. I. and Kammash, T. (1994). HYDRODYNAMIC FUEL CONTAINMENT IN AN OPEN-CYCLE GAS CORE NUCLEAR ROCKET. In *11th Symposium on Space Nuclear Power Systems*. American Institute of Physics. [16](#)
- [41] Poston, D. I., Kapernick, R. J., Dixon, D. D., Amiri, B. W., and Marcille, T. F. (2008). Reference reactor module for the affordable fission surface power system. In *AIP Conference Proceedings*, volume 969, pages 277–284. [61](#)
- [42] Reid, R. S. (2004). Alkali Metal Heat Pipe Life Issues. Technical report, Marshall Space Flight Center, National Aeronautics and Space Administration. [25](#)
- [43] Robert Bromberg, P., Romie, F. E., and Verdes Es, P. (1967). United States Patent Office METHOD OF AND MEANS FOR INCREASING THE HEAT TRANSFER CAPABILITY OF A HEAT. Technical report, TRW, Inc. [61](#)
- [44] Schmidt, G. (1996). SYSTEM TOPAZ INTERNATIONAL PROGRAM. Technical report, Space Technology Directorate AIR FORCE MATERIEL COMMAND KIRTLAND AIR FORCE BASE, NM 87117-5776. [11](#)
- [45] Shure, L. I. and Schwartz, H. J. (2020). Survey of Electric Power Plants for Space Applications. Technical report, NASA. [80](#)

- [46] Stokely, C. and Stansbery, E. (2006). Identification of a debris cloud from the nuclear powered SNAPSHOT satellite with Haystack radar measurements. Technical report, Barrios Technology. [10](#)
- [47] Trump, D. J. (2019). Presidential Memorandum on Launch of Spacecraft Containing Space Nuclear Systems — The White House. [7](#)
- [48] Truscello, V. (1983). SP-100, the US Space Nuclear Reactor Power Program. Technical information report. Technical report, Jet Propulsion Laboratory. [10](#)
- [49] Turner, M. J. L. (2010). Rocket and Spacecraft Propulsion: Principles, Practice and New Developments. [26](#)
- [50] Vijra, N. and Singh, P. (2015). An Experimental Study of Thermal Performance of Concentric Annular Heat Pipe. *American International Journal of Research in Science*, pages 15–179. [59](#)
- [51] Voss, S. S. (1984). SNAP Reactor Over. Technical report, Air Force Weapons Laboratory. [10](#)

Appendix

Appendix A

Nuclear Thermal Rocket Methods Development

A Prior Art

Beginning in 2006, a joint effort was undertaken between Los Alamos National Laboratory and NASA's Marshall Space Flight Center to develop an analytical capability for the study of NERVA-derived nuclear thermal rocket systems. Over a period of approximately two years, a design process was developed that consisted of three components: (1) a RocETS engine system model, (2) a spreadsheet-based thermal and mechanical design tool, TMSS-NTP, and (3) a FORTRAN code that generates MCNP models with input from TMSS-NTP. The project ended rather abruptly, and this capability sat unused from approximately 2007 until 2011, during which time Rick Kapernick and Ben Amiri both left Los Alamos National Laboratory. Unfortunately, the tool set was not readily usable as the NCPS project began. TMSS-NTP was tied to an Excel version and third-party solver which are no longer supported, and could not be run to completion. Both parts of the nuclear tool set were designed with NERVA derivatives in mind, and as a consequence, the overall nuclear process puts heavy emphasis on the design and performance of tie-tube elements. The RocETS models developed by NASA-MSFC were lost in the intervening years, though some effort was undertaken to recreate that modeling capability.

A description of each of the components, and the information exchange between them follows, as envisioned and implemented in the 2006 timeframe.

RocETS Model

RocETS, or Rocket Engine Transient Simulator, is a NASA-Marshall Space Flight Center developed tool used extensively in the design and analysis of rocket engines. Not surprisingly, given that every rocket flown by NASA to date has relied on chemical reactions for power, RocETS was limited as applied to nuclear rocket systems. NASA undertook the task of adapting RocETS to support steady-state engine calculations incorporating a simple “black box” reactor model, one whose parameters would be determined by TMSS-NTP and NTRgen.

Thermal-Mechanical Spreadsheet for NTP (TMSS-NTP)

TMSS-NTP is an evolution of a number of thermal and mechanical calculation tools developed at Los Alamos National Laboratory, primarily focused in this case on aligning the nuclear design requirements with those of the rocket engine system. TMSS-NTP takes the core inlet and outlet conditions from RocETS, along with a set of system performance metrics (e.g. hydrogen supply conditions for pump operation, nozzle cooling). A thermal analysis is performed for an average and peak fuel element, and the tie-tube components are sized to meet turbopump requirements. Initially, these calculations are performed using assumed power shapes; they are updated with MCNP tallies when the first iteration is completed. A parameter sweep is performed, changing numerous parameters that influence the neutronic performance, including total fuel mass and length-to-diameter ratio. Curve fits are generated, and passed to NTRgen. Typically, a run of TMSS-NTP would take about 8 hours.

Nuclear Thermal Rocket Generator (NTRgen)

NTRgen is a tool for generating MCNP models of nuclear thermal rocket engines.⁽¹⁾ With a relatively simple engine description, NTRgen creates models for a variety of reactor conditions: cold shutdown, normal operating temperature (including geometric changes due




Figure A.1: Prior Nuclear Rocket Design Process Overview

to thermal expansion), and immersion accidents (quartz sand, wet sand, and water). NTRgen also contains an optimization routine that, in conjunction with the thermal performance curves from TMSS-NTP, works to meet performance criteria while minimizing system mass. NTRgen produces 3-D power deposition contours that are passed to TMSS-NTR.⁽¹⁾ Process Overview Figure A.1 presents the overall flow of information in this design process. Given a set of mission requirements, a RocETS model of the system can be developed and analyzed. From the RocETS model, the core inlet and exit pressure and temperature, and tie-tube inlet and exit pressure and temperature drive the majority of the TMSS process. The combination of fuel-to-tie tube ratio and tie tube power requirement set the flow and power of each tie tube.

With a power distribution from MCNP, TMSS locates the peak power fuel element and calculates the average axial power shape. It is possible, though not likely in normal operation, that the peak fuel temperature will not be located in the fuel element with the highest power.




Figure A.2: TMSS Fuel Element Analysis

(Exit-peaked axial shapes in lower power elements will have the potential to see higher peak fuel temperatures.)

Using an “average” fuel element and a “peak power” fuel element, TMSS calculates coolant pressure profiles and temperatures in coolant and solid materials. Varying the number of fuel elements and fuel element dimensions over predefined ranges, “acceptable” core designs from a thermal performance standpoint are generated. An “acceptable” design must meet coolant temperature, pressure and flow rate requirements from RocETS, satisfy maximum coolant temperature and coolant temperature rise limits, and was envisioned to include some stress and fatigue limits (not implemented). A set of acceptable designs results in curves as seen in the following figure, whereby particular parameters can be varied by NTRgen while maintaining acceptable thermal performance. In this case, as an example, if NTRgen needed more reflector worth, it could increase core height. Changing the dimensions along TMSS curves will result in an increased fuel hex dimension, which at first might seem counterintuitive if fuel volume is being preserved. However, because the fuel elements are




Figure A.3: TMSS Tie Tube Analysis

longer, the coolant channels must be enlarged to maintain pressure drop, which turns out to be the larger effect in this case.

The thermal analysis for a tie tube module turns out to be much more complex than that for a fuel element. The tie tube serves two functions for the core itself, providing a moderator material (ZrH) and structural support for the fuel elements. Both the moderator and the structural tubes have temperature limits much lower than the fuel, and must be cooled. The hydrogen that cools the tie tubes is then used to drive a turbopump. TMSS, working with a set of tie tube moderator powers from MCNP, solves for dimensions of each of the tie tube components, attempting to maximize the amount of moderator while meeting material temperature and stress limits and meeting pressure drop constraints imposed by RocETS. To accommodate the effect of conduction from the adjacent fuel elements, a 1-dimensional (cylindrical) approximation is made of the fuel element, surrounding the tie tube block. A tie tube design is generated for each reactor configuration passed to NTRgen.

As previously discussed, NTRgen takes parametric curves from TMSS and attempts to arrive at a configuration that, already meeting T-H requirements, meets the nuclear




Figure A.4: TMSS Zoning from 2006 LANL-MSFC Effort

design criteria. These include sufficient reactivity to operate at temperatures corresponding to full power operation, sufficient control drum worth to shut down a cold reactor, and will eventually include a requirement for systems maintaining subcriticality during a set of launch accidents. Because of the large amount of void space within the core, these designs will be unable to remain subcritical without in-core poison material either in the form of control rods (which would displace fuel) or fine wires inserted in the coolant holes that would be withdrawn in-space prior to first use. NTRgen is also responsible for the generation of nodal power distributions, and implements the enrichment zoning method to confirm its performance (see figure A.4). Typically, when TMSS is run with a reasonable power distribution, a second TMSS run is sufficient to confirm the performance of the design generated by NTRgen. Very limited iteration has been performed between the TMSS/NTRgen processes and RocETS. This process works reasonably well for steady-state “point designs” useful for informing broad decision-making regarding technology selection.




Figure A.5: Flow Rate Solver

While it is by far the most complete and comprehensive nuclear rocket design tool to-date, it does not give any consideration to operation at anything other than full power, does not consider startup and shutdown hydrogen flow, transient behavior, or the effects of azimuthally non-uniform power distributions (as in the case of a malfunctioning control drum).

B Integrated Thermal and Neutronic Calculations in NTRgen

One of the most significant changes to the process is the integration of thermal calculations into NTRgen. Previously, the data exchange between NTRgen and TMSS was cumbersome, and TMSS runtimes were significant. Part of the computational expense for TMSS was the need to provide thermal solutions over a wide range of parameters, since the nuclear design was not known.




Figure A.6: Coolant Hole Diameter Calculation

The first step in this process is implementation of a model to calculate fuel element thermal performance. Unlike the prior work, the method here solves for the flow rate in a particular fuel element, given its dimensions and inlet and outlet conditions. Since the core inlet and outlet conditions (namely pressure and temperature) will be determined by the engine system, this seems reasonable. It provides for a relatively simple fuel element calculation, that can be repeated for each individual fuel element independent of the rest of the reactor (once MCNP-derived power distributions are available).

This change suggests an opportunity to introduce a new approach to balancing the core outlet temperatures (again, maximizing ISP/minimizing peak fuel temperature). Where prior analysis had considered enrichment zoning, changing the power produced per unit volume fuel, it is now possible to consider zoning based on flow channel diameter. There are

three first-order effects here, all working in the same direction: fuel volume (and hence power) decreases with increasing coolant hole diameter, flow increases (lowering exit temperature), and temperature rise within the fuel decreases. *** This will become much more significant once a transition to LEU fuel is introduced, where reductions in enrichment will be much more of a burden on system mass. From a design standpoint, the process for developing a low-enriched nuclear rocket concept is substantially similar to that for an HEU system. However, to keep the mass penalty for low-enriched fuel to a minimum, it is not practical to further reduce enrichment to balance fuel element power. Instead, fuel geometry must be modified. As previously discussed, the method for doing this will be to develop an algorithm to search for coolant hole diameters to balance the power and flow requirements for each fuel element. The complexity of this process is not overly significant for these systems, as burnup is sufficiently low that a beginning-of-life analysis is sufficient.

Early testing of the thermal models shown suggests a performance improvement from approximately 1 minute to less than 1 second for a single fuel element with 10 axial nodes. Following this, the module was modified to utilize an arbitrary nodalization for the fluid and heat transfer calculations, independent of the coarse (typically 10 axial node) MCNP discretization.

B.1 Channel Flow Rate Calculation

As shown in figure A.5, the flow rate calculation is an iterative process that starts with an initial guess of flow rate, temperature, and pressure profiles. The initial guess assumes a linear pressure drop, and temperature rises proportional to the fractional power in a node, where N_{ax} is the number of axial columns and T_i and P_i are the bulk temperature and average pressure at node i :

$$P_i = P_{in} - (i - 1) * (P_{in} - P_{out}) / N_{ax} \quad (\text{A.1})$$

$$T_i = T_{i-1} + (T_{out} - T_{in}) * \frac{PF_i}{N_{ax}} \quad (\text{A.2})$$

Next, the linear heat rate at each nodal location is calculated using the total enthalpy rise and the nodal power peaking factor, where L is the channel length:

$$q'_i = \dot{m} (h(P_{i+1}, T_{i+1}) - h(P_i, T_i)) * \frac{PF_i}{L/N_{ax}} \quad (\text{A.3})$$

Nodal Fuel Temperature and Thermal Expansion

From here, a call is made to a function that calculates the average and maximum temperatures at a particular nodal location, and performs a one-dimensional thermal expansion calculation. This requires calculation of a number of familiar fluid mechanics quantities (Re , Pr , Nu), and the use of the shape factor correlations described in section C.

$$h_c = \frac{Nu * k_c(P, T)}{d - 2 * t_{coating}} \quad (\text{A.4})$$

where d is the diameter of the fuel hole, $t_{coating}$ is the coating thickness, and $k_c(P, T)$ is the thermal conductivity of the hydrogen coolant, at the nodal pressure and temperature.

Temperature rise in the coating is an ordinary convective heat transfer and cylindrical heat conduction calculation:

$$T_{coatID} = T_i + q' / (\pi * d * h_c) \quad (\text{A.5})$$

$$T_{coatOD} = T_{coatID} + q' \cdot \log \frac{d}{d - 2 * t_{coating}} \frac{1}{(2\pi * k_{coating})} \quad (\text{A.6})$$

T_{coatOD} is the temperature at the inner diameter of the fuel meat, and is the reference temperature used along with the shape factors to calculate peak and average nodal fuel temperatures:

$$T_{fuel,avg} = T_{coatOD} + \frac{q'}{SF(P, d) \cdot k_{fuel}} \quad (\text{A.7})$$

$$T_{fuel,max} = T_{coatOD} + \frac{q'}{shape.factor_{max}(P, d) \cdot k_{fuel}} \quad (\text{A.8})$$

$$d = d_{cold} + d_{cold} \cdot CTE(T_{fuel,avg}) \cdot (T_{fuel,avg} - 293) \quad (\text{A.9})$$

Both average and maximum fuel temperatures are returned to the calling function, along with the expanded diameter.

Resuming Flow Rate Calculation

With a thermally expanded fuel node, we can calculate the pressure drop through the node.

$$v_{head} = \frac{1}{2}\rho \cdot \left(\frac{4\dot{m}}{\pi \cdot d^2} \right) \quad (\text{A.10})$$

$$\Delta P = v_{head} \left(\frac{2\rho_i}{\rho_{i+1} - 1} - (1 - (d_i/d_{i+1})^2) \cdot \left(\frac{2\rho_i}{\rho_i + \rho_{i+1}} + ((f_i + f_{i+1}) \cdot \frac{L/N_{axcol}}{d_i + d_{i+1}} \cdot 2 \frac{\rho_i}{\rho_i + \rho_{i+1}}) \right) \right) \quad (\text{A.11})$$

At the exit, the pressure is compared to the chamber pressure, and the mass flow rate guess is adjusted until the fractional difference between them is less than the convergence criteria.

B.2 Shutdown Model

Like any nuclear reactor, stopping the sustained chain reaction in a nuclear thermal rocket does not stop the production of thermal energy from delayed fission and radioactive decay. While the operating times of the reactor are very short, reducing the decay power, the high power density and poor passive heat rejection properties of the reactor system in space make decay heat removal a challenge.

Most nuclear thermal rocket development efforts have not reached a stage where shutdown cooling (i.e. heat removal after the first burn, during non-operating periods, for the duration of the mission) is seriously considered. The author developed a series of models to estimate (and hopefully bound) the amount of propellant required for this heat removal.

The first calculation was simply to estimate peak fuel temperatures as a function of decay power given only radiative cooling to space through the rocket nozzle throat. As one would expect, this results in extreme temperatures. While a simple calculation, this obviates the need to actively cool the system.

In the next model, an attempt was made to quantify the lower bound of propellant inventory required. If one assumes that propellant could be perfectly introduced into each channel, gaining heat energy from the fuel as it passes through, and exit at the mixed-mean operating temperature of the reactor, we can calculate the amount of propellant required to deliver that scenario as a lower bound. This is essentially an adiabatic lump-heating calculation in a quasi-steady state. At each time interval, the decay power is calculated, and the flow rate required (integrating the change in enthalpy over the pressure and temperature change from inlet to exit) is determined.

For the mission profile used in the SCCTE program, this results in a lower-bound estimate of 241 kg of propellant, essentially just to keep the reactor from heating further than just after operation. This is a small, but non-negligible amount, on the order of 1% of the total for the mission. Unfortunately, it would be virtually impossible to maintain minimum flow rates necessary across the large number of channels at these low total flow rates. It would also be difficult to know the film drop with any reasonable confidence. Increasing the flow rate to a level where the pressure drop through the fuel elements was significantly large relative to plenum and exit losses would require an unacceptable propellant inventory.

Next, a model considering a pulsed cooling method was developed. Instead of constant flow, this supposes an adiabatic heating period with no coolant flow, lasting until a peak fuel element reaches a temperature threshold. At that time, full-flow cooling is initiated, continuing until the peak element temperature falls below a lower threshold. This model exposes perhaps one of the bigger challenges of nuclear thermal rocket engine operation: changes in coolant flow result in a very rapid change in fuel element temperatures. A full-flow pulse well less than a second in duration cools the system more than 1000K.

A more advanced model attempted to quantify the point at which heat removal through the tie tube structures would maintain a reasonable peak fuel temperature. This model shows that tie tube cooling would be sufficient once decay heat falls below 0.5% of rated power (see Figure A.7. This would result in a significant reduction in coolant flow required, as tie tube flow can be modulated and the number of channels is significantly lower than the number of coolant channels (typically 19 or more coolant holes per element, 2 to 6 elements per tie tube).




Figure A.7: Fuel temperatures with only radiative cooling as a function of decay power. Note that decay power falls to 1% very quickly after shutdown.

Examining each of these approaches suggests a combination of methods provides the best result, using full-flow cooling briefly after shutdown, then transitioning to tie-tube cooling.

C Shape-Factors Correlating Finite-Difference Analysis to 1-D Heat Transfer Calculations

This component is currently in use within NTRgen. The following figures show a summary of results from a series of 2-D fine-mesh finite difference calculations performed using Thermal Desktop and SINDA. Calculated temperature profiles are shown in figure A.8. The nodal temperatures and associated areas were then used to identify maximum and volume-averaged temperature rises. The automated meshing routine places higher node densities near boundaries and geometric detail, so each computational node has a different area, with that area used to weight the nodal temperatures in calculating the average. The analysis performed to generate the chart below was handed off for implementation in the




Figure A.8: Thermal Desktop Shape Factor Calculations

TMSS calculations of peak fuel temperature. Correlations were generated there between pitch-to-diameter ratio temperature rise resulting in the following(22):

$$SF_{avg} = \exp \left[0.1421 \cdot \left(\frac{P}{D} \right)^4 - 1.560 \cdot \left(\frac{P}{D} \right)^3 + 6.502 \cdot \left(\frac{P}{D} \right)^2 - 12.67 \cdot \left(\frac{P}{D} \right) + 12.44 \right] \quad (\text{A.12})$$

$$SF_{max} = \exp \left[0.05931 \cdot \left(\frac{P}{D} \right)^4 - 0.6892 \cdot \left(\frac{P}{D} \right)^3 + 3.101 \cdot \left(\frac{P}{D} \right)^2 - 6.752 \cdot \left(\frac{P}{D} \right) + 8.208 \right] \quad (\text{A.13})$$

such that $T = T_{ID} + \frac{q'}{SF(P,D) \cdot k}$ estimates the temperature of interest, either the peak or average temperature within the truncated triangular region. Note that the shape factor is dimensionless, and does not depend on any particular unit of measurement.

These correlations allow reasonably accurate steady-state temperatures, and have been used within the new thermal/flow solver routines in NTRgen, without any significant computational expense.

Vita

David Dixon obtained Bachelor of Science degrees in Nuclear Engineering and Electrical Engineering from North Carolina State University in 2006 and 2007. In 2009, he was awarded a Master of Science degree in Nuclear Engineering at the University of Tennessee. David's first work in the nuclear industry began in 2002 as a student, working at Global Nuclear Fuels (GE Energy) in Wilmington, North Carolina. From there, David went on to work as a student at Los Alamos National Laboratory in Los Alamos, New Mexico, beginning during the final year of Project Prometheus. It was at Los Alamos that David was introduced to the world of space nuclear systems.

Upon completion of his Master of Science degree, David returned to Los Alamos full-time, immediately swept away from the intended course of work to aid in the response to the Deepwater Horizon oil spill. David was part of a team that was attempting to characterize the magnitude of the spill in support of government litigation against the rig operator. As part of this process, David adapted a two-phase water/steam fluid flow model he had constructed as an undergraduate for a boiling water reactor to model the three phase fluid system present in the oil well.

Shortly thereafter, David was again called on to aid in the Department of Energy's response to the Fukushima Daiichi nuclear accident. In the process, David called upon his experience in the waterjet industry to devise an experiment to prove that an abrasive waterjet could be used to bore a hole in 1-meter thick concrete with rebar more than an inch in diameter, without an ignition source. Having been exposed to the waterjet industry as a child and having worked in the family waterjet business, David attempted to located an industry expert to speak to a process for doing so, calling on a number of industry experts and university professors. Finding none, and in fact finding those experts generally didn't

think such a feat was possible, David set out to demonstrate his approach. Calling on friends and family, a similar reinforced concrete, with similar strength characteristics was poured and cured over a few days, and within a week of the first discussions, traveled to North Carolina to build a CNC machine with the capability to execute the test, write the algorithm to perform the cutting, and complete the boring process. In a three-day weekend, all of those things were completed. It should come as no surprise in reading this dissertation that a nuclear reactor test, with a small budget and rapid timeline, originated with him.

Not long after the DUFF experiment, David worked with NASA's Marshall Space Flight Center for a short time on a nuclear thermal propulsion system design effort, before starting his own business which developed a line of products for the automotive industry that have significantly improved the safety of nighttime driving, with notable reduction in fatalities when vehicles are equipped with his products.

David is a native of Ohio, has lived all over the country as a child and young adult, and since 2006 has traveled the country in a large motorhome, with southwestern South Dakota as his permanent home.