“Evolvable” Systems are the Key to Successful Space Reactor Development

A White Paper by David Poston – Lead reactor designer for Kilopower/KRUSTY, and designer for numerous special-purpose reactor projects. This is solely my personal opinion: email spacenukes@gmail.com.

Summary of Key Points

- The most enabling/impressive space reactor capabilities are at high power-density and temperature; the promise of these systems continues to bring agencies back to pursue space reactors despite 50+ years of time/money wasted.
- The hardest part of space reactor power system development does NOT necessarily lie in demonstrating technology (fuel, structure, pump, power conversion, etc.); for most concepts, it lies in demonstrating the performance and control of the integrated system.
- Simple, predictable system performance is the key to successful development because of the high cost/risk of ground testing, inability to create/benchmark complex models, and lack of capability for in-space reactor instrumentation and control.
- The path to high-performance space reactors (Nuclear Thermal Propulsion (NTP), human nuclear electric propulsion (NEP), industrial mining, and colonies on Moon/Mars) is to develop “evolvable” systems; which start simple and then use prior system tests/experience to ensure successful prediction/verification of each subsequent generation.
- Kilopower was created as an evolvable space reactor concept, which NASA leadership supported until just after the successful KRUSTY test in 2018. Kilopower is evolvable because subsequent generations can maintain all or most aspects of the predictable, self-regulating power system dynamics and control that were proven by the KRUSTY test. Initial Kilopower generations will require little or no technology development.
- Developing “evolvable” systems is more important than developing “scalable” technologies.

INTRODUCTION

There is one question every space reactor designer will hear more than any other: “What is the maximum power and/or lowest mass your reactor technology can provide?” This might be a good question if there were space reactor models on the shelf or at least a few potential options that had been actually tested. Instead, I believe the question itself is largely at the root of our past 50+ years of failed programs; i.e., people tend to focus more on the future of space reactors than the present. Note: in this document, space reactor is used as vernacular for space fission power system (SFPS). There is always more interest/funding for advanced paper concepts and/or technologies that could hypothetically provide more in the future (because the latter sells programs). Recently, the problem has gotten worse because we have no longer have qualified SMEs (Subject Matter Experts) in space reactor development; e.g., people who have the experience to inform decision-makers of what is real and what is hypothetical (if not blatantly false). For established engineering applications (e.g., proposing a new rocket-engine concept) there are plenty of qualified SMEs that can call B.S. or identify shenanigans in the proposal, unfortunately, this capability does not exist for space reactors.

The question that people should actually ask a space reactor designer is “What can you point to that provides confidence that your concept will work as advertised?” This would include claims for component technologies, mass, cost, schedule, system performance, control and operation, safety, etc. Another informed question might be "What is unique about your approach that will allow it to succeed, as compared to the previous dozens of failed programs?" Unfortunately, these sorts of questions are rarely asked. The lack of space reactor SMEs creates an environment that rewards those who simply tell the customer what they want to hear (as opposed to those who understand and convey true engineering reality), especially if they work for institutions that claim to have expertise in designing/developing new reactors (which is by definition false advertising, because a novel reactor concept hasn't been designed, developed, and operated in over 50 years – except for KRUSTY[1]).

WHAT HAS BEEN ACCOMPLISHED IN THE PAST (PRIOR STEPS)?

Almost Nothing. We've spent billions of dollars on space reactor programs over the past 60 years. We don't have any reactor capability in space, nor have we ever had a successful test of any new space reactor concept in 50+ years (again, except for KRUSTY). See Table 1 for a list of previous programs (green is power, blue is propulsion).

Table 1: Abbreviated List and Cost Estimates of US Space Reactor Programs.

Decade	Project	Estimated Cost	Cost Today	Reactors Tested	Flight Units
50s-60s	SNAP	~$380M	$2.40B	~6	1
50s-70s	Rover/NERVA	~$2B	$12.00B	~20	0
70s-80s	SPAR	~$10M	$0.06B	0	0
80s-90s	SP-100	~$1B	$2.50B	0	0
80s-90s	MMW program	~$50M	$0.13B	0	0
1990s	NEBA (bimodal)	~$5M	$0.01B	0	0
1990s	Topaz (US effort)	~$50M	$0.09B	0	0
1990s	SNTP/Timberwind	~$200M	$0.34B	0	0
2000s	Affordable Reactor Prog.	~$5M	$0.01B	0	0
2000s	JIMO	~$400M	$0.53B	0	0
10s-20s	FSP	~$40M	$0.04B	0	0
10s-20s	NCPS/NTP	~$200M	$0.21B	0	0
2010s	Kilopower	$18M	$0.02B	1	0
Total			**~$18.4B**	**~27**	**1**

Every program in Table 1 (and numerous others) started with the belief they would deploy a reactor, but only SNAP did. There are many reasons that cause large government programs to fail, but one aspect of space reactor programs today is unique, <u>a lack of capability and expertise</u>. All terrestrial reactors that have been deployed to date have been based on greater than 100 reactor tests in the 50s/60s. This experience, knowledge, and capability are gone, because of 50+ years with no new reactor concepts tested (until KRUSTY). This lack of “real” expertise continues to lead to the proposal and adoption of paper reactor concepts that have little chance of success without significant ground testing (true for both space and terrestrial programs). This is essentially a chicken-and-egg problem: a) we can’t develop capability/infrastructure with knowledge/experience, and b) we can’t develop knowledge/experience without capability/infrastructure. The way to resolve this dilemma is to take small, simple steps, which can evolve to highly capable, enabling systems.

WHAT COULD WE DO NOW (FIRST STEPS)?

There are two classes of mission that can be successfully completed today if a focused program pursues a reactor power system that allows “integrated simplicity”. Integrated simplicity is an approach to chart the easiest path to deployment, combining simplicity in design, technology, procurement, fabrication, assembly, transport, operation, instrumentation, control, testing, safety, and the necessary government and regulatory approvals. All facets of development and deployment need to be recognized and incorporated into any proposed space reactor system/program.

<u>In-space power option for the first step</u>: Low-mass space power and electric propulsion from 1 to 10s of kWe. Our ability to study the outer planets and their Moons is limited. State-of-the-art deep-space spacecraft not only have limited electric power (<<1 kWe), but more importantly have no practical ability to orbit far-off worlds. This class of system would allow years of study at close range (versus very short fly-bys), and high power for increased diagnostics and signal bandwidth. Other uses for this type of system would likely emerge if it was available.

<u>Surface power option for the first step</u>: Robust 10s of kWe for early surface power missions: Estimates of the ideal power system size for early Moon/Mars mission range from 10 to 30 kWe (depending on the level of redundancy, flexibility, transportability desired in the architecture). Power would be used for life-support, science instruments, infrastructure, and in-situ resource processing. Even after multi-megawatt reactors are deployed for large settlements, the transportability of these systems will allow expansion to remote regions and charging stations along the routes between settlements.

There are three key criteria for any potential first step concept or program to succeed: the concept must 1) have a simple and robust safety approach, 2) use technologies that exist, are procurable, and have been demonstrated to operate in the intended environment, and 3) be able to operate successfully in space without the need for a ground nuclear-powered system test. The latter criteria are needed because a nuclear-powered system test will almost certainly have too much programmatic/cost uncertainty for a successful first step. The key to eliminating the need for a ground nuclear test is to design a concept that a) has direct heritage to a previous reactor test, or b) has simple and predictable physics, reactor dynamics, and system controllability, and allows prototypic non-nuclear (e.g. electrically-heated) testing. The system with the biggest head-start in this direction, and also meets all of the above is Kilopower[2], but it is acknowledged that others could come up with concepts that could achieve the same goal.

WHAT COULD WE DO LATER (FUTURE STEPS)?

The easy part of this debate, and what perpetually keeps the dream of space reactors alive, is the type of space fission power systems that we could have in the future. Fission offers orders of magnitude better performance than any other space power source. Much of the benefit of SFP is tied to the very high energy density (J/kg) potential of fissile material, which can allow high powers and/or long lifetime at a relatively low mass. In addition, there is no fundamental limit to the power density that fission can provide; the power density is limited by engineering, not physics. All other near-term power sources have fundamental physical limits: solar and radioisotope power are limited by power density, while chemical and batteries are limited by energy density. Space fission power systems also offer great flexibility in terms of being able to operate in diverse environments (in particular as compared to solar), and the ability to provide robust electrical power from kilowatts to megawatts.

Figure 1 contains two plots that provide a highly-generalized view of potential space fission power performance. The plotted metric is the specific power, or Watts of electricity per kilogram, which is usually the most important measure of performance. These estimates of specific power include all aspects of the power system: e.g. reactor, shield, power conversion, heat rejection, etc. Both charts contain the same data but are split into separate Kilowatt and Megawatt scales to allow more resolution. These kinds of charts are far too simplistic to inform design or programmatic decisions, but they are meant to portray the simple, intuitive idea that pursuing higher performance involves higher risk.

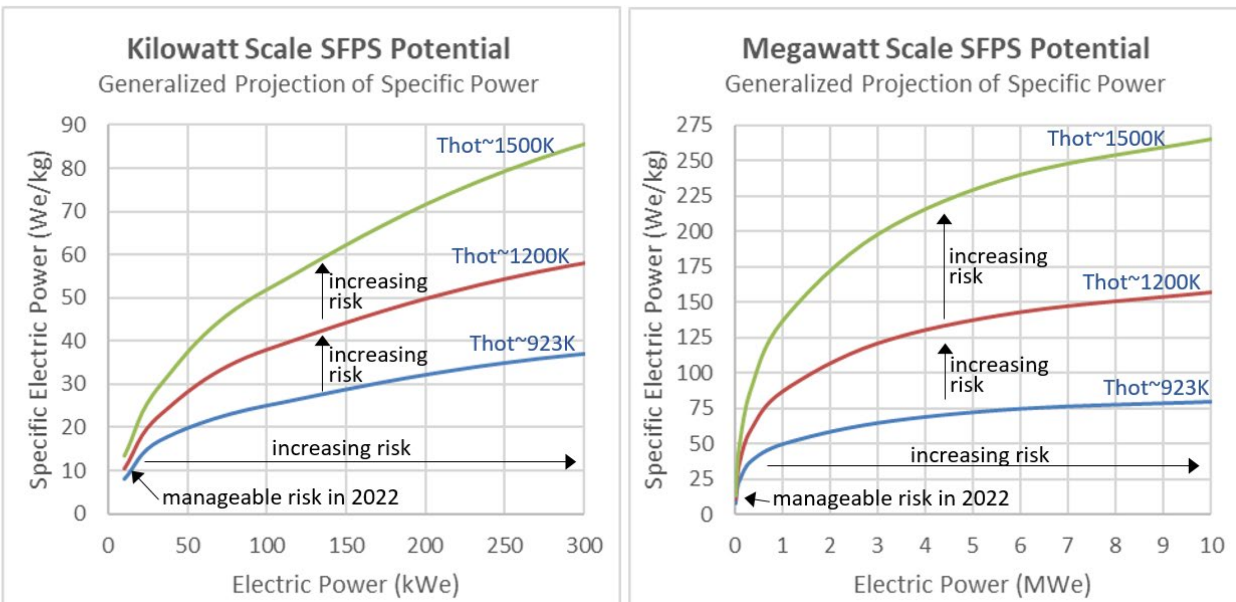

Figure 1: Simplified depiction of space reactor specific power versus risk.

Moving to higher powers (left-to-right on Figure 1) increases specific power because the reactor and power conversion (depending on technology) can produce significantly higher powers with modest increases in mass. As power grows, the curves flatten as the reactor becomes less criticality limited (i.e. when engineering issues other the neutronic criticality limit the performance), and at powers >MWe the heat rejection mass (~linear with power) becomes a bigger fraction of the mass. Moving to higher temperatures (upward on Figure 1) improves performance by either increasing power conversion Carnot efficiency and/or reducing the size of the radiator, recuperator, etc. Again, these are highly generalized charts. Changes in reactor, power conversion, and radiator technologies can significantly change specific power, but the overall trend will remain intact.

The high-level message of Figure 1 is that there is a strong correlation between performance and risk. (Note, in this context risk represents the probability of success of system development, deployment, and utilization). Manageable risk represents the “what we could do now” discussion above. Moving forward from there, risk increases with both power and T_{hot} (the power conversion hot-end temperature). The lowest risk curve (T_{hot}=923K) is based on design calculations for Kilopower-type reactors that use High Assay Low Enriched Uranium (HALEU). Stirling power conversion is assumed at <100 kWe, and then Brayton conversion at higher powers.

Risk of higher power: Moving to higher powers (left-to-right on Figure 1) increases risk because of higher fuel burnup (swelling and fission gas), increased irradiation damage to materials and components, higher thermal gradients and peak temperatures, increased stresses and strains, more difficult component thermal management (e.g. cooling of reflector and shield), and dozens of additional issues. Increased power can also make safety and control requirements significantly harder to meet, because of increased core size and coolant area (potential flooding).

Risk of higher temperature: Moving to higher temperatures or T_{hot} (upward on Figure 1) increases risk because it requires both a higher temperature reactor AND a higher temperature power conversion system. The risks imposed by higher temperature are generally harder to overcome than those imposed by higher power because of the need to qualify high temperature materials. Development programs for high temperature materials and components is often very lengthy and risky, especially given the unique requirements of a space reactor (nuclear effects, launch loads, long-lifetime, lack of maintenance, etc.).

Other performance risks: The most important performance metric not included above is system lifetime, which impacts numerous engineering risks. In some respects, lifetime increases risk in the same manner as power (e.g., irradiation, burnup), while other lifetime risks are more correlated with temperature (e.g., creep, chemical interactions, diffusion). Overall, systems that use hydrated moderators are much more sensitive to lifetime than unmoderated systems. Other performance metrics may depend more on the specific application; (e.g., reliability and radiator size).

HISTORY SHOWS THAT SYSTEM EVOLUTION IS THE PATH TO TRULY ENABLING SYSTEMS

Almost every space technology in use today has arrived at advanced performance via evolution: solar power, battery storage, control systems, etc. Today, NASA and DoD are both lucky beneficiaries of technology evolution – the Falcon-9. From Day 1, the goal of SpaceX has been to develop rockets to take humans to Mars. To their credit, SpaceX was smart enough to know that they did not have the experience or resources needed to develop Starship, the Falcon-Heavy, or Falcon-9. They knew they needed to take manageable steps. Falcon-1 was the perfect first step because it was doable (just barely) within the resources they had, and it was just capable enough to have some utility. Of course, the great value of Falcon 1 was not the engineered system itself, but the experience the company gained and the subsequent support that ultimately led to Falcon-9 and F9-Heavy, both of which have been game-changers to NASA and DoD capability.

Some may say that “lucky” in the paragraph above is unwarranted because NASA/DoD put a lot of money into SpaceX – I think those people are incorrect. They were lucky because SpaceX was willing and able to fund the necessary first step (Falcon-1) that led to these visionary and enabling technologies. Just as SpaceX did not have the experience or resources to go straight to Falcon9/Starship, the USA does not currently have the experience or resources (capabilities and funding) to develop an advanced game-changing space reactor. We may have had what was required in the ‘60s and ‘70s when the US successfully developed dozens of new/novel reactor concepts, but that capability no longer exists.

One of the reasons that NASA/DoD appear averse to funding a first-step in fission power is likely the same reason that they were not significantly interested in funding something like Falcon-1; the performance does/did not knock their socks off. Our ability to develop a space reactor with "wow-factor" performance right out of the box has been essentially eliminated as we've lost capability, AND because the bar has been continuously raised by the establishment with the evolution of solar panels and other energy/storage technology. Unfortunately, the current prevailing thought in the US is that we do not need to start simple, and we can go straight to high-performance space reactor systems. I think this thought goes unchallenged because of our lack of SMEs with space reactor capability/expertise. And regardless of my opinion, this prevailing thought can be shown to be incorrect simply by examining the past five decades of failed programs (and in all of those cases decision-makers felt confident that they were smart enough to go straight to advanced reactors as well). <u>If we ever want to unleash the potential of fission power in space, we will either need a white-knight investor or a visionary leader in NASA/DoD to commit to a small, doable first step.</u>

WHAT MAKES A SPACE REACTOR SYSTEM "EVOLVABLE"?

Evolvable can be viewed as a "system-level" term, as opposed to scalable or extensible, which can be viewed as "technology-level" terms. An evolvable system not only begins with a system that can be deployed and operated with manageable risk but also can incrementally pursue improved follow-on concepts with manageable risk. Depending on the lessons learned from the first successful deployment, a manageable second step might move either to the right (higher power) and/or upward (higher temperature) in Figure 1.

The short answer to why "system evolvability" is more important than "technology scalability" is rooted in the oft-used axiom – "if it doesn't fly, it doesn't scale", where "fly" in this case implies not only deployment but successful operation. The hardest part of space reactor power system development does NOT necessarily lie in demonstrating technology (fuel, structure, pump, power conversion, etc.); for most concepts, it lies in deploying the system and demonstrating the performance and control of the integrated system. Simple, predictable system performance is the key to successful development because of the high cost/risk of ground testing, inability to create/benchmark complex models, and lack of capability for in-space reactor instrumentation and control.

COMMON MISTAKES OF DECISION-MAKERS OVER THE YEARS

<u>Decision Maker Mistake #1</u>: "If we're going to spend >$1B on a space reactor, it better have the performance we want". This argument might be valid if every concept would cost billions to develop, but a simple low-cost system can allow you to complete the first flight reactor for as low as ~$100M. The success of a first flight reactor gives you the capability and experience to succeed on the next evolution to an advanced high-performance system. This first space reactor flight informs you on which advanced technology/system to pursue, on the best safety approach, etc. and this makes it is more likely the next evolved system will be successful, and as a result significantly reduces the cost of what would have been a likely-to fail >$1B program.

<u>Decision Maker Mistake #2</u>: "The first space reactor we deploy must have fuel/technology that's scalable to very high performance." This argument could be valid if you needed to spend a lot of time and money developing ANY technology to deploy the first system; however, if the technology for the first system already exists, then it is better to deploy the simple system first (thus learning how to design, develop,

deploy and operate a space reactor) and work on the advanced technology for the second system in parallel (and if a specific technology development fails, you now have the first evolutionary step to build upon).

Decision Maker Mistake #3: "We're not going to waste time and money taking a small first step when the majority of space nuclear experts tell us we can go straight to high performance." This mistake is rooted in the aforementioned lack of SMEs with capability and expertise in space reactors resulting from 40+ years of failed programs. Compared to other space technologies (rockets, solar) there are no experts in "how to design/develop a new reactor"; thus, anyone can say they are an expert and if they put on a good show people will believe them (especially if they tell the customer what they want to hear). See Table 1 for a list of failed programs as validation of this statement.

KILOPOWER: EXAMPLE OF AN EVOLVABLE SYSTEM

The Kilopower reactor was designed, by intent, to be evolvable. The primary goal was to meet the aforementioned criteria for a successful 1st step: simple/robust safety, existing/procurable technology, and simple/predictable system operation. On top of this, the design offered a simple and achievable confirmatory testing approach. The Demonstration Using Flattop Fissions (DUFF) test confirmed some aspects of the simple reactor physics[3], but more importantly, that nuclear-powered testing can be affordable for a prototypic reactor. The Kilopower Reactor Using Stirling TechnologY (KRUSTY) test[4] then demonstrated the simple, fault-tolerant, and self-regulating operation[5] of a flight-prototypic Kilopower reactor. KRUSTY used highly enriched uranium (HEU), and while very low-power Kilopower reactors benefit from HEU, systems designed to provide >10 kWe use HALEU fuel, which allows for higher power (by limiting fuel swelling) and fortunately does not significantly change the dynamics and operation of the power system.

Kilopower has clearly checked the box as a concept that could achieve a successful first step (i.e., successful operation in space) and is currently the only concept with a test that confirms this. The remaining question is how Kilopower could evolve to higher performance systems. First, it is important to recognize that it is impossible to deterministically chart the path of evolution because evolution builds upon knowledge and experience gained from previous generations (i.e., things that are not known now). However, it is also important to have identified at least a few high-performance evolutionary options that appear feasible. Three potential technology changes can move the current KRUSTY-based Kilopower design to a substantially higher power (moving left-to-right on Figure 1). 1) A change from UMo fuel to a fuel that is better suited for higher burnup and temperature (likely UO2, or perhaps UN). 2) A change from a monolithic block of fuel to a monolithic block of structural material (e.g., stainless-steel, super-alloy, or refractory metal), in which rods of fuel are placed. 3) A change from Stirling converters to a Brayton power conversion system. There might also be a transition from a heat-pipe-cooled to a gas-cooled core depending on how the manufacturing and integration of high-power heat-pipe reactors proceeds. The order and timing of all of these changes will depend on how these technologies perform in prior generations, as well as component technology development. For example, a switch from Stirling to Brayton would depend on how well the Stirling converters perform in early generation(s) and how difficult higher-power integration appears to be, as well as how well low-power Brayton technology and control evolves. Moving upward on Figure 1 would similarly depend on how technology development proceeds for higher temperature core and power conversion components. Ideally, the three envisioned technology changes would occur in different evolutionary steps. The mantra of legendary NASA program manager

John Casani is that “a program should never try to tackle more than one significant technology change at a time”.

The final, and perhaps most important feature of Kilopower evolvability is the ability to preserve dynamic and operational heritage to previous generations. This feature eliminates the need for nuclear system testing (and potentially numerous design-build-test iterations) to develop a working concept. The 1st-generation system must be able to rely on the confirmatory operational tests of KRUSTY, the 2nd-generation must be able to rely on the 1st, and so on. Many possible Kilopower evolution paths will allow this progression to take place.[6]

CONCLUSION

The path to high-performance space reactors (NTP, human NEP, industrial asteroid mining, colonies on Moon/Mars, etc.) is to develop “evolvable” systems; which start simple and then use prior system tests/experience to ensure successful prediction/verification of each subsequent generation. Developing “evolvable” systems is far more important than developing “scalable” technologies. This distinction is not a subtle/esoteric difference, it is at the core of why space reactor development is difficult and different from traditional system engineering (largely due to lack of expertise and the ability for nuclear system testing). Programs over the past 40 years have emphasized enabling performance over development risk, and all of them failed because they did not make sufficient progress within the touted cost and schedule. This issue was recognized decades ago by Admiral Hyman Rickover[6], in a letter he penned in 1953 concerning “academic versus practical reactors”, and more recently in a white paper concerning “real versus paper” reactors[7]. The ironic consequence of this cycle of failure is that if any previous decision-maker had decided to go the simpler path in the past, we could already be on our 2nd or 3rd generation of space reactors, which would probably have even better performance than the advanced technology we are pursuing today. Too many decision-makers are concerned about what they consider “dead-end” technologies; (i.e., there is a specific aspect of an achievable, real reactor that will not get them to where they ultimately want to go). The real dead-end is never making enough progress to learn and build anything. Any technology that leads to a working space reactor creates forward progress towards any highly advanced systems.

References:

1) Poston, D., 2021. “Space Fission Power: Kilopower and KRUSTY”. In: Encyclopedia of Nuclear Energy, vol. 3. Elsevier, 2021.
2) McClure, P., Poston, D., Gibson, M., Mason, L, Robinson, R., “Kilopower Project: The KRUSTY Fission Power Experiment and Potential Missions”, Nuclear Technology 206:sup1 2020.
3) Poston, D., McClure, P, Dixon, D., Gibson, M., Mason, L., “Experimental Demonstration of a Heat Pipe-Stirling Engine Nuclear Reactor”. Nuclear Technology 188 (3), 2014
4) Poston, D., Gibson, M., Godfroy, T., McClure P., “KRUSTY Reactor Design”, Nuclear Technology, Vol. 206:sup1, 2020.
5) Poston, D., Gibson M., McClure, P., “Scalability and Evolution of the Kilopower Reactor Power System”, to be published NETS 2022.
6) Rickover, H., Letters to editor. Nature 243, 312, 1973
7) Poston, D., “The Persistent and Growing Gap Between “Real” and “Paper” Special-Purpose Reactors”, https://www.spacenukes.com/technical-papers, 2020.