

The Persistent and Growing Gap Between “Real” and “Paper” Special-Purpose Reactors

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

Today’s special-purpose reactor (SPR) programs in the US (space fission power, NTP, and Microreactors) are rapidly moving away from potentially real, achievable reactor concepts. The increasing trend in the US has been to build interest and funding in SPRs by promoting high-performing paper reactor concepts that have no ties to real reactor experience. In all cases, there are varying degrees of development risk with respect to unproven materials, fabrication options, procurement, launch safety, etc. These risks are daunting for many of the systems being pursued today, but most people can understand these risks to some degree. However, there is one flaw in today’s current space programs that most people cannot grasp – they are pursuing concepts that have complex and uncertain reactor dynamics; i.e. how the reactor power system will operate (e.g. startup, state points, transients, failure response, stability, nuclear burnup effects, etc.). Concepts of this type will almost certainly require numerous design, build, and test cycles to get the system to work adequately. Pursuing these concepts is a serious mistake because the US has no ability or will to perform multiple repeated nuclear ground tests, and unsurprisingly no intent to spend tens of billions of dollars to create that ability. Many programs have identified this testing dilemma, and have proclaimed they will have a reactor “demonstration” instead of a test program, but they have almost no idea what the chances are that the demonstration will work. History and logic both indicate that a new reactor with complex reactor dynamics will not succeed the first time it is operated. A regulator will be hesitant to approve testing of this type of reactor as well. Unfortunately, this risk will not appear until the end of the program, after years have been wasted; even worse, the cycle of failed programs will continue.

The development of SPRs, especially space fission power systems, is unique from any other advanced technology. Testing of new nuclear systems has become almost prohibitively cumbersome and expensive – this is surely part of the reason there has not been a successful SPR deployed for many decades. In addition, the testing of systems that operate in space presents other unique challenges. The combination of these two, nuclear and space, puts space reactors in an unfortunate class of their own. For a traditionally engineered system, the development approach can allow significant uncertainty in system performance, and a vigorous test program can be used flush out performance issues and ultimately optimize performance (even for terrestrial reactors, as long as there is a chance to perform exploratory low-power reactor operations, and make subsequent changes, before extended full-power operation). This allows the design and development process to focus on what most view as traditional engineering: materials, procurement, fabrication, component development, and assembly. Alternatively, due to the cost and risk of testing, a space fission system must be designed so that a system “test” is in effect an actual “demonstration”, with as few uncertainties and unknowns as practically possible.

The primary problem that has led to a cycle of failed programs is the inability of decision-makers (and too often the PowerPoint engineers that inform them) to balance benefit versus risk. For example, the majority of NTP performance risk could be eliminated by using HEU fuel vs LEU fuel, or by allowing a ~2000 or 3000 kg mass increase. Someone would be correct to point out the use of HEU is a very large program

risk, or that 2000 kg is a major performance hit¹, but that same person probably has no idea how huge the risk is of assuming that a complex NTP system will work adequately without numerous, perhaps dozens of ground tests. For an LEU surface power system, a realistic moderated system has the potential to save perhaps 200 kg over a much simpler KRUSTY-derived system, but the number of development risks associated with the moderated system make it costlier and lengthier to develop, and more importantly, have a substantially lower chance of ultimate success. An externally-moderated LEU Microreactor might save a factor of two in fuel mass/cost, but has a substantially higher probability of not working, if it ever advances far enough through technology development to be built in the first place. In this example, a simple to develop/operate reactor that uses twice as much fuel may not be commercially viable, but it would be more likely to save lives in logistical military or disaster scenarios (because it would be far more likely to actually get deployed). Switching to unmoderated systems won't come close to solving all of the issues for any of these applications, but can give them a fighting chance.

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 reactors, which would probably have even much better performance than what they shooting for now. 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 reactor power system creates forward progress towards highly advanced systems. Plus if an actual system is developed, some users may find that it is beneficial as is, albeit not as enabling as what they wished for.

Reactors with complex and uncertain dynamics (coupled neutronic, thermal-hydraulic, and mechanical performance) are not impossible, perhaps best evidenced by the existence of Boiling Water Reactors. Even the US' two successful SPRs (SNAP-10A and naval reactors) were moderated systems. However, all of these programs had several nuclear tests prior to deployment, which is not what today's programs are assuming. Unlike SNAP-10 and naval reactors, the only NTR functional prototype ever tested (i.e., one that used a prototypic configuration of reactor, turbopump, nozzle, etc.) was dynamically very simple. The NERVA XE' system used simple 1-pass flow, used an open-cycle turbopump, used HEU, did not use a moderator, and operated at much lower temperature and pressure than any NTR concept proposed today (plus other dynamic simplifications like lower reactivity feedback coefficients, lower hydrogen worth, lower power density, lack of reliance on thermal insulation, etc.). These simplifications allowed the successful demonstration of an NTR with a rated Isp of only 710 seconds, but it was a real reactor, and they were evolving towards a solution. Many people tout that over 800 seconds was demonstrated, but that was only a simple core test, not an NTR test; others have been saying for decades that NTR startup and operation has been proven, but that was only for this highly simplified system, which has virtually no dynamic similarity to even the SNRE concept that was proposed at the end of the NERVA program. Overall, XE' had 6 major reactor dynamic simplifications over any NTR concept that is being pursued today! The proposed complexity of today's concepts is especially hard to believe given that even the ROVER/NERVA engineers didn't progress far enough to begin flight-system design, after 10 years and 10 billion dollars (2020 equiv.); and they lived in an environment where nuclear tests could be iteratively conducted about

¹ The mass advantage of moderated systems is only significant at lower thrust, where concepts are criticality limited (as opposed to being limited by power density, stresses, reactivity swings, etc.), but at 75k lbf (the target thrust for NERVA) the advantage disappears.

every 6 months (20 reactors in total were tested). This disconnect in paper versus reality is truly astounding in today's NTP programs.

While the focus of this white paper is focused on the underappreciated risk of reactor dynamics, there are of course many other differences between paper reactors and real reactors. The functionality of reactor instrumentation and control is usually taken for granted in paper reactor designs. Non-deployable terrestrial reactors have 50+ years of reactor experience that has led to functional instrumentation and control solutions (often via lots of redundancy). An SPR will not have the luxury of virtually unlimited mass and space to deploy sensors, and will see environments and temperatures different from any existing reactor (especially if deployed in the vacuum of space). In-core temperature measurements may not be possible at all, which would present a major problem for any system that has complex reactivity feedback and sub-cooled regions (i.e. more than one significant temperature coefficient). The desire for low-mass might also make it difficult to ensure enough shielding to avoid control single-event-upsets and other radiation induced issues (a typical reactor will use generous amounts of concrete and dirt to shield controls and make this a non-issue). SPR instrumentation and control also must survive the potential vibrations and impact associated with deployment and operation. All of the above will create a system where it is hard to know which instrument signals can be relied upon to command the control mechanisms, as well as the difficulty of ensuring control element movement in these unique environments. The difficulty of obtaining reliable instrumentation and control is another compelling reason to avoid concepts with complex and uncertain reactor dynamics. Best case is a concept that does not need any real-time or automated reactor control, something akin to the Kilopower reactor concepts.

Another risk that is often ignored by paper reactor concepts is nuts and bolts engineering/structure. The more complex and unproven a system is, the more that the design will have to be altered to accommodate reality (likely with significant mass increase). There were over a dozen design changes in KRUSTY (an extremely simple reactor) as the mechanical design progressed, and prototyping of components and assemblies was performed. More than a dozen other changes were caused by requirements that were not evident when the project initiated, or by delivered parts that did not meet spec., etc. Fortunately, KRUSTY was simple enough and had ample design margin to fully succeed, and no major issues arose during fabrication or testing. Alternatively, the far more challenging and complex NRX test series of the NERVA program had 16 "significant structural anomalies" that occurred over the 6 ground test design iterations, including cracking, flaring, rubbing, breaking of various components.

The current trend away from real reactors is especially disheartening because it was only 2 years ago that the first truly new reactor, KRUSTY, was tested in the US in 40+ years. Prior to that, the US (DOE, DOD, NASA) had spent billions of dollars over 40 years pushing paper reactor concepts, without a single system completed, and amazingly not even a single nuclear-powered test! Except for perhaps some changes in naval reactors over the years, which non-coincidentally have been successfully executed by the same agency as the recent KRUSTY test – the NNSA. Naval Reactors is the only US success story in SPRs, and this is largely attributable to the wisdom and perseverance of Admiral Hyman Rickover. Most reactor engineers are familiar with Admiral Rickover's famous words on academic vs practical reactors in 1953. Unfortunately very few heed his advice, even though Rickover hands-down made the greatest contribution to SPRs in history.

"An academic reactor or reactor plant almost always has the following basic characteristics: (1) It is simple. (2) It is small. (3) It is cheap (4) It is light. (5) It can be built very quickly. (6) It is

very flexible in purpose ('omnibus reactor'). (7) Very little development is required. It will use mostly off-the-shelf components. (8) The reactor is in the study phase. It is not being built now.

The academic-reactor designer is a dilettante. He has not had to assume any real responsibility in connection with his projects. He is free to luxuriate in elegant ideas, the practical shortcomings of which can be relegated to the category of 'mere technical details.'

Unfortunately for those who must make far reaching decisions without the benefit of an intimate knowledge of reactor technology and unfortunately for the interested public, it is much easier to get the academic side of an issue than the practical side. For a large part those involved with the academic reactors have more inclination and time to present their ideas in reports and orally to those who will listen. Since they are innocently unaware of the real but hidden difficulties of their plans, they speak with great facility and confidence."

The last paragraph perfectly describes what has happened during the past 40+ years of special purpose reactors in the US; but even after 40 years of failure, decision-makers are still wooed by the promises of academic/paper reactor designers. Every SPR project since Rickover's retirement has not only failed miserably, but has not even reached the point of an actual, "real" reactor test – until KRUSTY in 2018. Although I've heard there are some out there saying the KRUSTY wasn't a "real" reactor test – I would love to have that debate.

KRUSTY's success was rooted in "integrated simplicity". Integrated simplicity is an approach to minimize the integrated risk of successful development; including, materials, manufacturing, procurement, design, modeling, assembly, project management, safety, and regulatory approvals, and steady-state and dynamic performance. The reason KRUSTY succeeded was that there was a small tightly-knit interagency team that was able to collectively understand and make project decisions to navigate the simplest path to success. Much of today's ramped enthusiasm for SPRs is rooted in the success and positive press from the KRUSTY reactor test, which broke the cycle of failures through its simplicity. Enthusiasm in Microreactors was rooted in simple KRUSTY-like fast-spectrum systems as well. The ~1 MW Megapower concepts were generated from 2010 to 2012, and they were very well received in briefings to the DOE, DOD, and most importantly the Defense Science Board in the early-to-mid 2010s. The current program is moving away from simple concepts that would be much more likely be successful and thus save lives.

Finally, it was hoped that KRUSTY, through its unique success, would encourage others to break the endless cycle of programs that take on too much risk for a relatively small benefit (mass, unit cost, etc.). There is still time for today's SPR programs to shift to simpler concepts and succeed, as long as 1) decision-makers can find a way not to over-constrain requirements to prevent a real reactor from being built (and perhaps wait for the 2nd generation build to get truly enabling performance), 2) bureaucrats and managers can define success as project completion, instead of where the money flows, and 3) engineers can stop promoting paper reactors that have little chance of success (noting that I was one of those for many years myself – it's hard not to want please the customer and have fun while designing slick reactors to meet their wish list). Of course every program and concept is unique, so generalizations do not apply equally in all cases. I'm sorry if I offended anyone, but the importance of special purpose reactors to the US and humanity is too important for me to further suppress my opinions – although most everyone in the community already knows my opinion. For some reason, I don't think my message is getting relayed to the "right" people, whom I hope will find this document, consider what I have to say, perhaps read some of my technical papers and/or invite me to talk with them further.