

RESULTS OF THE KRUSTY NUCLEAR SYSTEM TEST

David I. Poston¹, Marc Gibson², Patrick McClure¹, Thomas Godfroy³, Rene Sanchez¹

¹Los Alamos National Laboratory, Los Alamos, NM 87545, 505-667-4336, poston@lanl.gov

²NASA Glenn Research Center, Cleveland, OH 44135

³Vantage Partners LLC, Brookpark, OH 44142

The Kilowatt Reactor Using Stirling Technology (KRUSTY) was a prototypic nuclear-powered test of a 5-kWt Kilopower space reactor. This paper presents results from the KRUSTY nuclear system test, which operated the reactor power system at various temperatures and power levels for 28 consecutive hours. The testing showed that the system operated as expected, and that the reactor is highly tolerant of possible failure conditions and transients. The key feature demonstrated was the ability of the reactor to load-follow the demand of the power conversion system. The thermal power of the test ranged from 1.5 to 5.0 kWt, with a fuel temperature up to 880 C. Each 80-We-rated Stirling engine produced ~90 We at a component efficiency of ~35% and an overall system efficiency of ~25%.

I. INTRODUCTION

The Kilowatt Reactor Using Stirling Technology (KRUSTY) was designed to be representative of a 5-kWt Kilopower¹ space reactor. KRUSTY was designed, developed, manufactured, and tested for <\$20M, with final testing completed in March 2018.

Kilopower reactors are designed, via simplicity, as load-following systems, i.e. the reactor will first-order provide whatever thermal power is demanded from it. The passive response of the reactor is similar to a household thermostat – if the temperature gets too cold the fission power/heat kicks on and vice-versa. The temperature or thermostat “set-point” is determined by the amount of reactivity insertion (i.e. the control rod position) and the passive temperature feedback of the reactor

KRUSTY was designed with flight prototypic materials and full-scale components to study the reactor dynamics at full power². The design allowed the testing of several nominal and off-nominal conditions, which ultimately verified that the reactor could tolerate any credible worst-case conditions without active control. A companion paper in these proceedings provides additional detail about the reactor, and describes the testing prior to the full nuclear system test³.

II. TEST OBJECTIVES

The overarching goal of the KRUSTY project was to show that a useful reactor power system could be designed, built, and tested quickly and affordably. Prior to KRUSTY, the biggest obstacle to space fission system development was the cost of previously failed programs. There were three high-level objectives for the KRUSTY nuclear

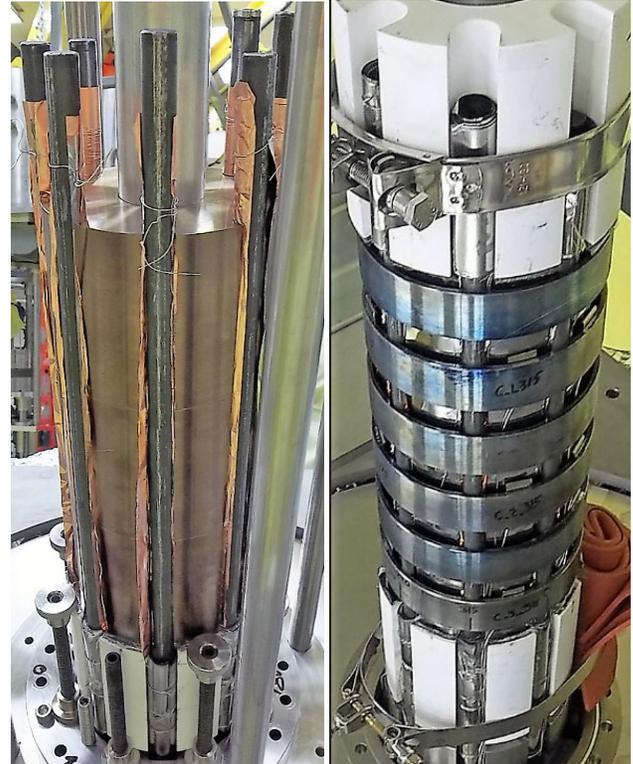


Fig 1. KRUSTY core assembly. Left photo shows heat pipes fit within the slots in the HEU UMo fuel. Right photo shows core after the installation of the Haynes 230 rings that clamp the heat pipes to the fuel (via interference fit). White BeO axial reflectors seen on top and bottom. Some parts in the photos are part of the temporary assembly fixture, which is later removed.

system test. 1) Operate the reactor at steady state with a thermal power output of 4 kWt at a temperature of 800° C. 2) Verify the stability and load following characteristics of the reactor during nominal and off-nominal conditions. 3) Benchmark codes and material data using with test results.

A test plan was devised that allowed as many reactor transients as possible, while also allowing enough settling time between transients to approach steady state. The test was limited to 28 hours to limit activation of the facility and to allow manageable staffing plan for operators, supervisors, and required facility and safety personnel. Following the 60-cent run³, the transient FRINK⁴ model was quickly benchmarked to give an informed prediction of the power/temperature oscillations. The updated FRINK model predicted an oscillation period of ~15 minutes (depending on the power draw) and that most transients



Fig 2. The KRUSTY “Power Conversion System” (PCS) consists of two 80 W-rated Stirling converters and six simulators (which mock the thermal behavior of the engines). At the bottom is the flange for the vacuum vessel. Over 100 thermocouple wires and connections are packed between the flange and the PCS.



Fig 3. The vacuum chamber (left) containing the core (inside lower left) and PCS (inside upper left) is lifted to be installed on COMET (right), which holds the large SS-304 shield (upper right) and the lower shielding on the platen (lower right)

would dampen in 3 or 4 oscillations. Thus, it was decided initiate a new transient every hour; except for a longer period at startup and the final, loss-of-active-cooling transients. It was very beneficial to have flexibility in the test plan, in order to incorporate knowledge gained on-the-fly during the criticals and the final test itself.

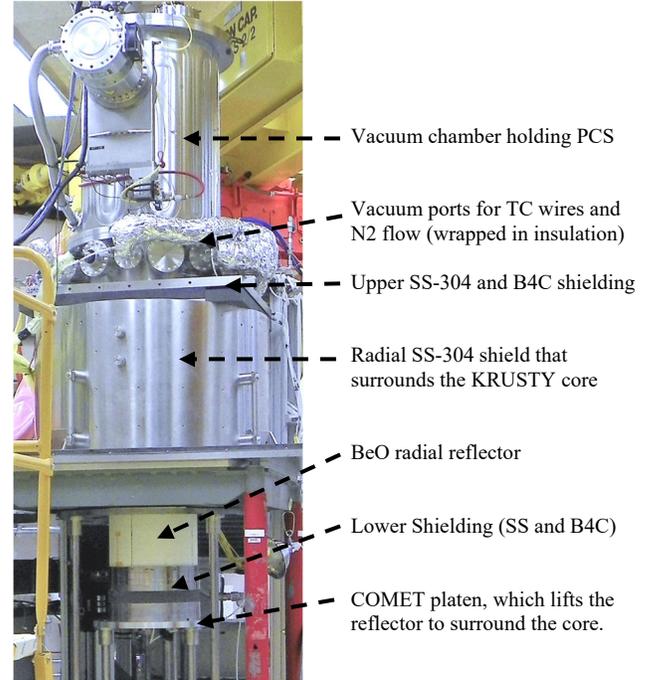


Fig 4. KRUSTY ready to go the morning of the final test.

III. NUCLEAR SYSTEM TEST

Figure 5 shows a condensed version of the results over the entire 28 hours of the nuclear system test.

III.A. What Exactly is Plotted?

Figure 5 and all subsequent plots display 3 types of measurements. The majority of the plotted curves depict fuel thermocouple (TC) data. These TCs are spring-loaded on the outside of the fuel; separated by a 0.001” layer of Mo to protect the fuel. A careful look at Fig. 1 reveals the edges of some rectangular spring clips and a few TC wires. During the warm criticals it was discovered that the coupling of the TCs to the fuel is poor at low temperature. However, once the core heats to a few hundred Celsius, core expansion appears to create good thermal contact, as evidenced by a <5 C variance between individual TCs. Each plot contains a red line that represents fission power, although the actual measurement is an Amp reading from a neutron detector, and fission power was later normalized to Amps³. The fission power is the total prompt energy (i.e. does not include decay power) distributed throughout the entire system: with ~94% to the fuel, ~2% to the reflectors and ~4% to the shielding. The amount of decay power depends on the prior fission history of the transient. A very rough approximation is that after ~1 minute of steady

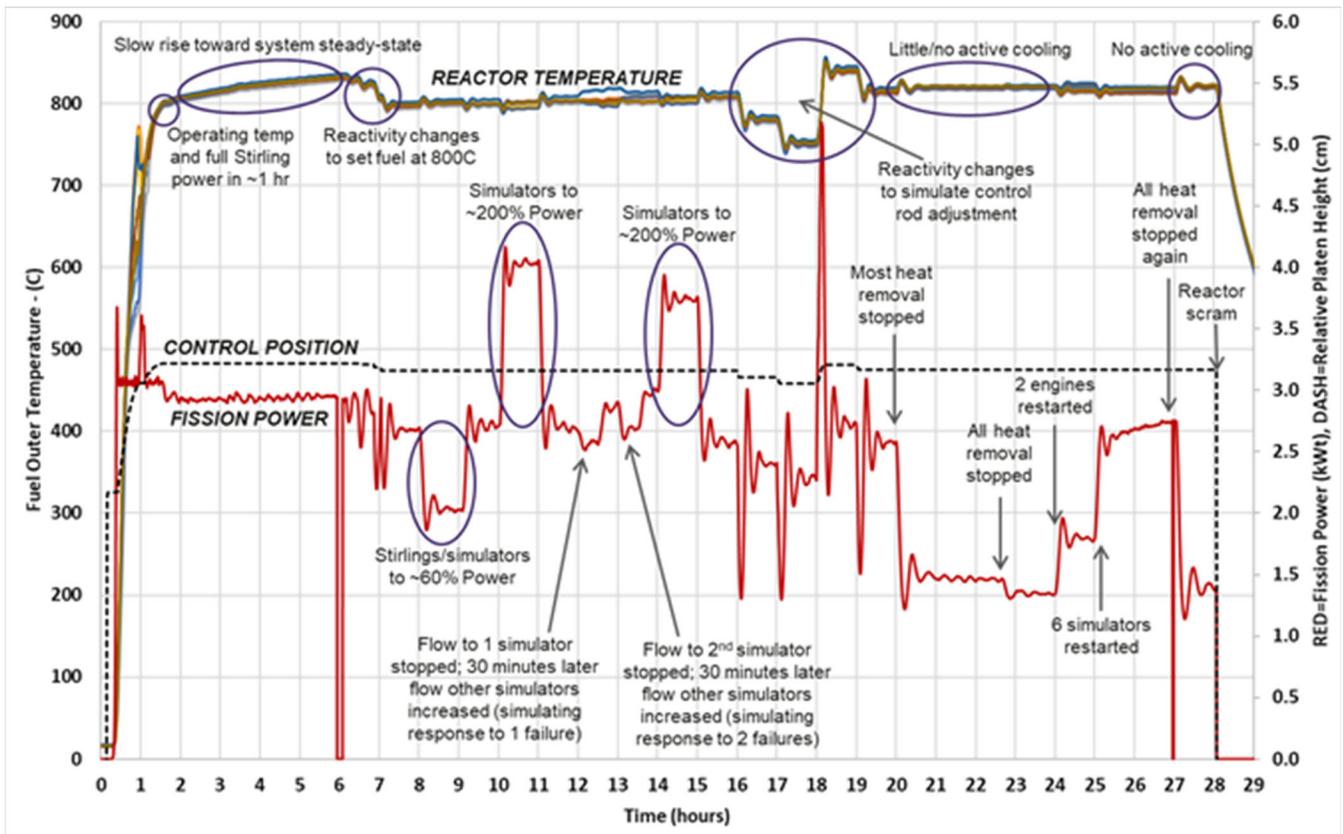


Fig 5. Power and temperature data from the full 28-hour KRUSTY nuclear system test.

operation decay power would be ~3% of fission power, then 4% after 5 min, 5% after 30 min, and 6% after 24 hours. So for most of the test, the actual thermal power is 5 to 6% higher than depicted by the red fission power curve.

A dashed black line on each plot depicts the relative platen position. KRUSTY's reactivity is changed by axial movement of the BeO radial reflector on the COMET platen (lift table). The value plotted is the distance from the top edge of the uppermost ring-clamp to the top edge of the BeO on the platen (somewhat arbitrary, but this allows the parameter to be co-plotted on the secondary axis with power, and may be easier to visualize).

To fit within page restrictions, the results in this paper are limited to those discussed above. Not included are TC readings for the heat pipes, Stirlings, simulators, reflector, shield, structure, and nitrogen (N₂) gas inlet/outlet, as well as vacuum pressure and the N₂ flow rate. Some of those measurements are needed to understand the more subtle changes in reactor power and temperature, and will be discussed in more detailed final KRUSTY documentation.

III.B. Startup

A closer look at the first 1.7 hours of the nuclear system test is shown in Fig 6. This test started the same way as the preceding warm criticals³ – with a 15-cent free run. The power increased exponentially until physical

heating created reactivity feedback and a subsequent power spike of ~3.7 kWt. Once the power dropped to ~3 kWt, reactivity was continuously inserted to keep power at ~3 kWt. Once the fuel reached 800 C, reactivity was no longer inserted, to allow testing of the coupled reactor/PCS dynamics. These activities, and all other major actions and phenomena are denoted with text and arrows on Fig 6.

The most important actions left off of Fig. 6 are the startup of the Stirling engines at T=1.12 hr, followed by the simulators at T=1.27 hr. The reactor was still being actively controlled at that time (via raising the platen when power dropped below 3 kWt), so there is no noticeable impact on power from these actions, nor a load following response.

III.C. Steady Power Coast and Reactivity Adjustment

After startup, the system operated nominally for several hours. The most notable feature on Fig. 5, from T=2 hr to T=7 hr is the slow fuel temperature rise. This was caused by 2nd order reactivity effects, most notably BeO heating and the axial expansion of the vacuum vessel extension. These effects are of significant academic interest, and will be fully addressed and benchmarked in the final KRUSTY documentation, but both are primarily caused by the limitations imposed by the ground test configuration (namely the existence of a vacuum vessel between the fuel and reflector). It is important that all 2nd order reactivity effects, whether for ground test or flight,

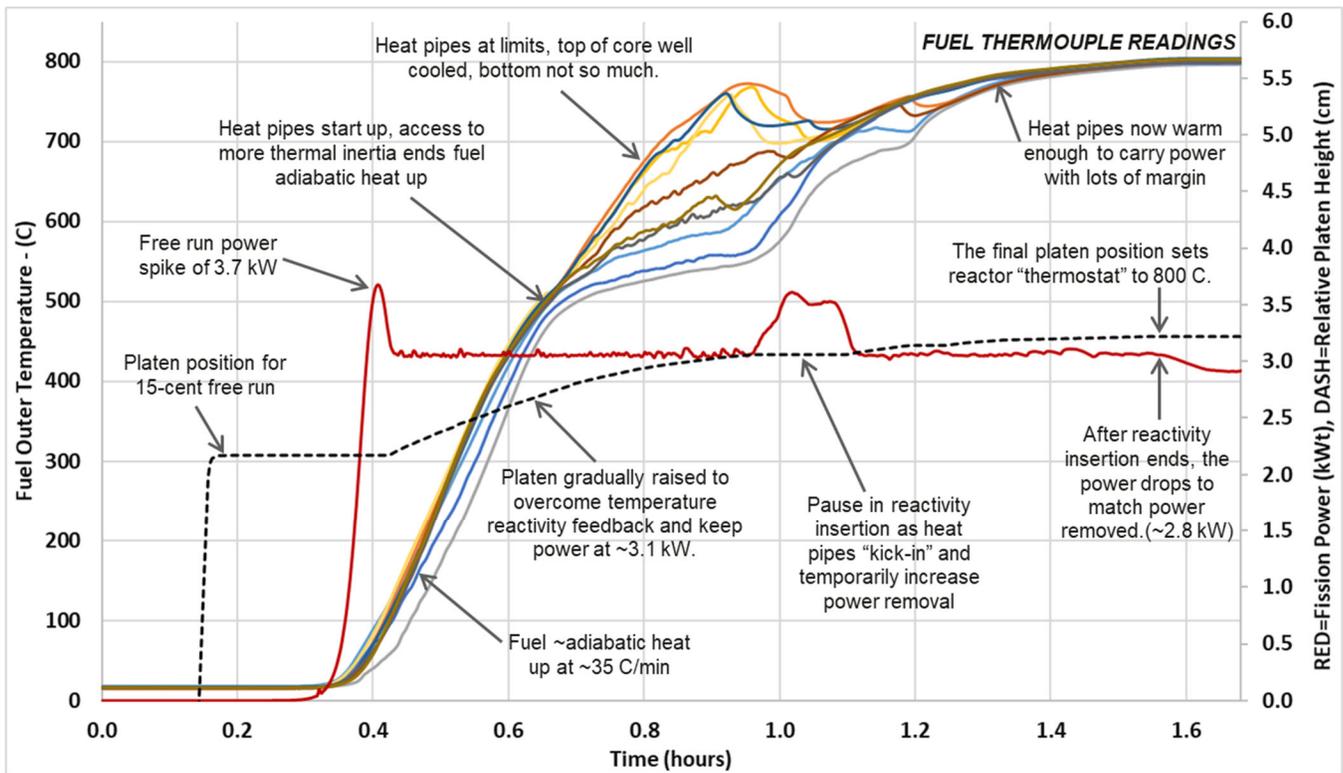


Fig 6. Startup data from the KRUSTY nuclear system test.

are very small relative to the fuel, just as the two effects occurring here. At $T=7$ hr the control (platen) was adjusted to bring the fuel temperature back to ~ 800 C, just as final tweaks might be performed remotely for a flight system.

III.D. Load Following Transients

Figure 7 takes a closer look at the load following transients from $T=8$ hr to $T=12$ hr. The test data confirms the expected load following characteristics. A decrease in heat removal causes a rise in temperature, thus a decrease in reactivity (via thermal expansion causing more neutron leakage) and thus a drop in power, which in turn causes temperature to fall back to the reactor “thermostat set-point” (via a slight overshoot and undershoot). The exact opposite occurs when power draw is increased.

A close look at the fuel temperatures on Fig. 7 shows that the fuel temperature is lower at higher power and vice-versa. This is expected because the reactivity set-point is based on the average fuel temperature (plus other 2nd order effects). Higher powers create a larger temperature gradient in the fuel, so the center runs hotter and the outer surface (the TC location) runs cooler. This is an important characteristic that needs to be considered when designing a load-following reactor – as thermal power to the PCS increases, the power is delivered at a lower temperature. This requires the ability of the PCS to generate more power at a lower efficiency. Design margins dictate the upper range of the load following ability, which ultimately represents the max-rated-power of the system.

III.E. Fault Tolerance Transients

Figure 8 plots the fuel TCs during the fault tolerance transients. At $T=12$ hr the flow was cut to the simulator at the 0° azimuth. As expected, the TCs near the 0° azimuth warmed up and the power dropped. At the same time, the TCs on the opposite side of the core (near the 180° azimuth) drifted lower, to maintain the thermostat set-point. At $T=12.5$ hr the N_2 flow to the other simulators was increased to simulate the power draw needed to return to nominal electrical power (i.e. if real Stirling converters were actually there to use it). At $T=13$ hr the flow was cut to the 180° simulator, and then at $T=13.5$ hr the flow to the remaining “working” simulators was again increased to simulate the return to nominal power. The thermal power at the new nominal-electric-power state point increases slightly to offset the slight drop in efficiency due to lower temperature heat pipes.

III.F. Reactivity Control Transients

Figure 9 shows the system response to reactivity change. At $T=16$ hr the platen drop of 0.5 mm causes a temperature drop of ~ 30 C. This decrease in temperature is roughly equal to the reactivity worth of the platen movement (0.5 mm at ~ 12 cents/mm = ~ 6 cents) divided by the reactivity temperature coefficient (~ 0.2 cents/C). The subsequent platen movements produced the same, expected behavior. The same thing would also occur if an internal control rod was moved. Zero-power criticals determined the control rod worth to be ~ 4 cents/mm near

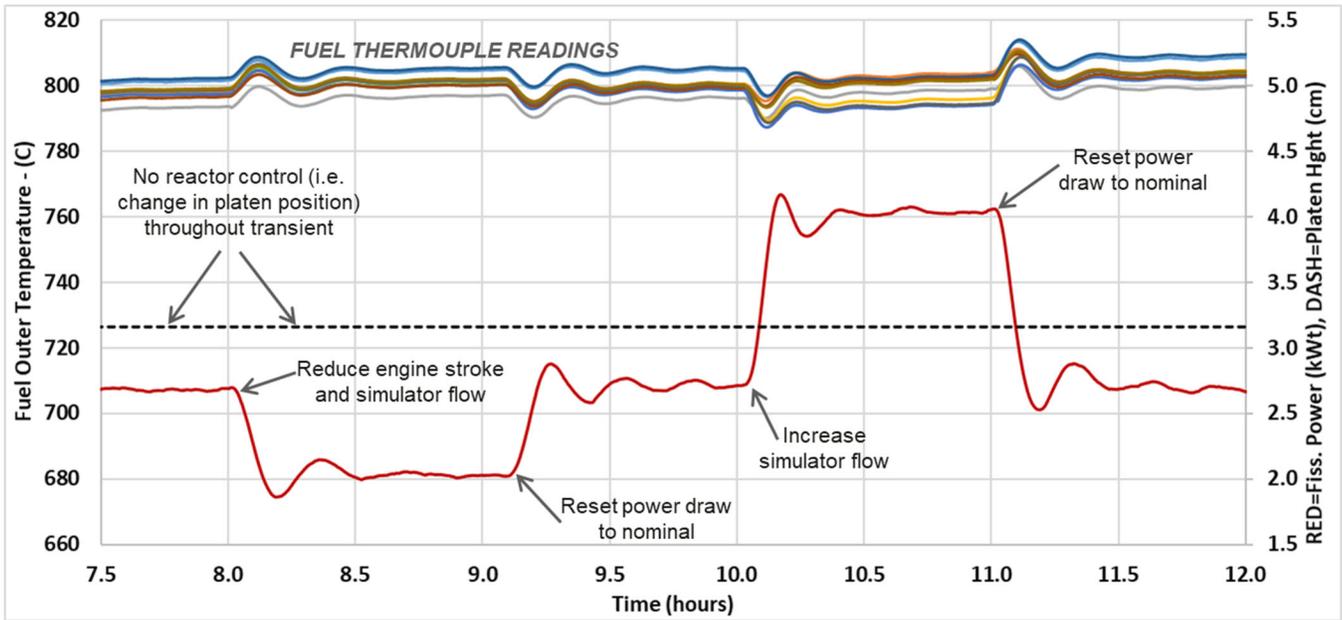


Fig 7. Load following transient data from the KRUSTY nuclear system test.

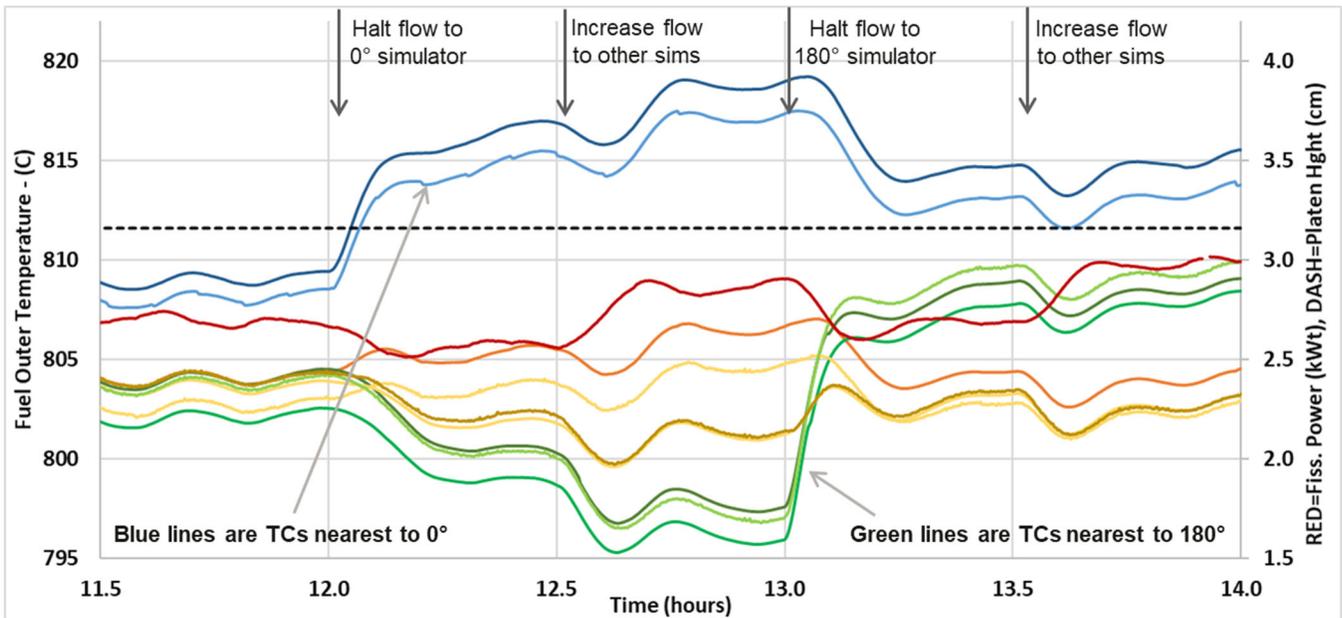


Fig 8. Fault tolerance transient data from the KRUSTY nuclear system test

the expected critical position³, thus it would need to be moved 1.5 mm to produce the same effect.

Another more subtle change on Fig. 9 occurs with the power level, which settles to a slightly higher level at increased temperature and vice-versa. This occurs because the converters and simulators draw more power when the core (heat pipes) are at a higher temperature and vice-versa.

VIII.G. Loss of Active Heat Removal

The final transient demonstrated the ability of the reactor to passively handle the loss of all active heat removal. As seen on Fig 5, at T=27 hr, the results of this

transient once again confirmed the ability of the reactor to passively load follow the core heat removal.

Unfortunately, the thermal simulator design and configuration was not ideal for this aspect of the demonstration. When the simulators were “on” (i.e. gas was flowing through them) the body of the simulator was kept very cool (near the inlet gas temperature of ~100 C). During nominal conditions, the power draw by the gas in each simulator was ~280 W and the radiative losses were estimated at ~40 W (thus ~320 W total drawn by the heat pipe). When simulator flow was cut-off, the temperature of the simulator and tubing increased dramatically (>500 C,

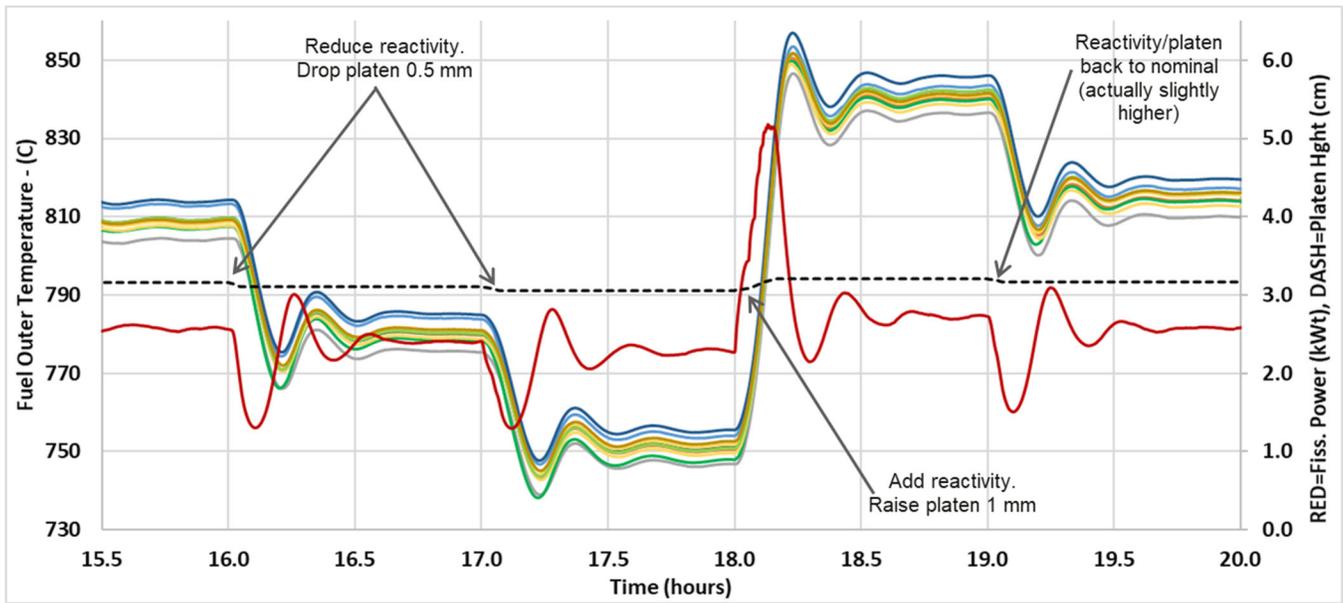


Fig 9. Control movement transient data from the KRUSTY nuclear system test.

as indicated by TCs in that region). In this no-flow condition the radiative losses increased to ~150 W per simulator, thus the total power draw was only cut by a about factor of 2. This is why the power at hour 27 on Fig. 5 only drops from ~2.8 kWt to ~1.4 kWt.

During this loss-of-flow transient the fuel temperature rose by only ~15 C before settling back to the ~800 C “thermostat” set-point. Pretest modeling predicted a ~30 C peak temperature rise⁵, which is consistent with ~15 C for a 50% loss. Note that the heat pipes feeding the actual Stirling convertors had a net power removal of ~320 W at nominal stroke, but only ~50 W when engine was shut down. So if the project budget had allowed 8 convertors, or if the simulator inlet region could have been better insulated, then the power drop caused by loss-of-active-heat-removal would have been more prototypic.

III.H. Electrical Output

The primary purpose of KRUSTY was not to produce electrical power; however, electrical power is space is indeed the goal of the overarching Kilopower project. Budget constraints limited KRUSTY to two 80-We rated Stirling convertors designed and built by Sunpower Inc. These convertors continuously produced ~87 We throughout most of the 28-hours, except for when stroke was reduced at T=8 hr and when engine was shutdown at T=20 hr. The engine thermal efficiency was ~33% (~50% of Carnot) depending on various parameters. The electrical power and efficiency changed noticeably as a function of reactor temperature during the reactivity control transients (T=16 hr to T=19 hr). During these transients, the average convertor power dropped to 78 We with fuel at ~750 C and reached 90 We with fuel at ~840 C. On average, a core temperature drop of ~30 C caused a ~4 We drop in convertor power (~1.5% drop in efficiency) and vice-versa.

Another important transient occurred at T=24 hr, when the Stirling convertors were restarted from zero-stroke (full shutdown). The engines had been sitting idle at very high temperature for 4 hours; thermally soaked to ~800 C, well above their nominal design temperature. The engines started immediately without a problem and provide a clean load-following response. As a bonus, at 800 C the engines provided a peak power of 105 We, and provided >100 We for more than a minute as the hot-end cooled down. Note that these convertors were undersized and are not the proposed flight convertors; they were used because of cost limitations. The flight convertors however are planned to use the same technology.

IV. HOW KRUSTY APPLIES TO FLIGHT KILOPOWER SYSTEMS

The biggest value of the KRUSTY test, besides showing that a flight-like system/technology worked, was to verify the simple and predictable dynamic response of a Kilopower power system – a heat-pipe-cooled, compact, fast-spectrum reactor coupled with Stirling convertors. There is very little thermal-neutronic difference between KRUSTY and any envisioned Kilpower system, from 1 kWe and 10 kWe, including HEU or LEU. In each of these possible systems, reactivity feedback (via fuel expansion) and heat transfer will be governed by the same physics and technology. Power level is surprisingly unimportant in determining how a neutron population behaves (because neutrons do not interact with each other), so extrapolating these reactor physics to higher power concepts is straightforward. What affects the neutrons is how power changes the temperature and geometry; therefore, nuclear dynamics will be similar as long as a reactor is engineered to provide similar thermal behavior to KRUSTY. The tolerance to failed components (heat pipes or convertors)

and the load-following power range might vary for specific designs; however, in most cases the design margins can be set to achieve the desired flexibility and reliability.

It is important to note that the difference between moving the BeO radial reflector (KRUSTY) and a B₄C internal rod (flight) is very small. The raising of the BeO increases reactivity by decreasing neutron leakage, while withdrawing the B₄C rod increases reactivity by decreasing neutron absorption. This difference will cause small 2nd order effects on power distribution and feedback, but as far as the neutron population (i.e. power) is concerned, there is very little difference between the two reactivity mechanisms, despite their different geometric locations. This is because KRUSTY is a very good example of a point-kinetic reactor, which occurs when the neutron mean-free-path is a significant fraction of the core geometry. In such a system, all regions of the reactor communicate very well with each other and regional feedbacks effects are negligible (except for how they affect integral reactivity). The reason that KRUSTY did not attempt to move the B₄C rod was because of the time required to gain safety qualification for the control mechanism. The actual mechanism will not be complicated; the only difficulty may be gaining confidence of long term operation in a radiation environment – something that KRUSTY would not have demonstrated regardless. In addition, low power/lifetime applications would not require B₄C rod movement after startup (because of low burnup reactivity loss), which could simplify qualification for that type of application. The other challenge with the control rod will be precluding inadvertent rod movement due to accident (particularly launch accidents) or human error.

Another KRUSTY-versus-flight difference is gravity: whether it's zero-g in space, micro-g during thrust, or gravity on any planetary body. The only significant way that gravity will effect operation is with heat-pipe performance. The KRUSTY heat pipes had a wick in the evaporator, but used thermosiphon action in the adiabatic and condenser regions. Kilopower heat pipes will have to be qualified to work in all possible environments – this is being addresses in the current Kilopower project. Note that the gravity effects will probably only matter during startup, when the heat pipes might try to operate at their limits; although a slow core heat-up should mitigate this issue. Once the heat pipes reach a temperature that provides substantial margin to all heat transfer limits, then there should be no problem. One thing made clear during KRUSTY testing was how well, and how quickly the 800 C heat pipes reacted to changes at either the core end (evaporator) or the PCS end (condenser).

Of course the flight PCS and heat rejection will be substantially different than KRUSTY, as well as shielding and various structural features. Launch and landing load analysis is now being incorporated into the Kilopower

design process. Also, lifetime effects were certainly not demonstrated by KRUSTY. From a nuclear perspective, the neutron fluences and fuel burnup are low enough that existing data provides high confidence for long life. Non-nuclear issues like mass diffusion between dissimilar materials or material creep need to be further examined, but can be mitigated with lower temperature (with drop in efficiency) if needed. A flight program will have to adequately resolve all of these issues and more.

V. CONCLUSIONS

KRUSTY demonstrated the operation and dynamics of a flight-like Kilopower system. As predicted, the reactor reliably load-followed the PCS power draw, and was able to accommodate potential system transients without a reactor control response. More so, the performance was very close to that predicted by the pre-test design and modeling tools, which gives even more confidence in designing and qualifying a robust flight concept.

The reactor core operated for over 24 hours at >800 C and >1 hr at >850 C. This is by no means a life test, but if final inspection is positive, it will provide added confidence for extended high temperature operation. Also, since the TC readings measure the outer fuel surface, the peak internal temperatures likely reached 880 C.

The reactor load-following ability was demonstrated from 1.5 kWt to 4 kWt during the test, and was only limited by the characteristics of the convertors and simulators. Likewise, the 60-cent critical³ demonstrated load following at 100 W of reactor power and below. Thus the entire testing suite effectively demonstrated full load following from 0 to 4 kWt of fission power. The reactor also operated at a fission power of >5 kWt for >5 minutes, but the PCS could not draw enough power to maintain steady-state.

Finally, the KRUSTY test was not only the first of its kind, but the first nuclear-powered operation of a truly new reactor technology in the US for over 40 years. The data and experience should help in all future reactor development efforts, and hopefully breathe life into the nuclear power community as a whole.

REFERENCES

1. M. A. GIBSON, et. al., "NASA's Kilopower Reactor Development ..." NASA/TM—2017-219467 (2017).
2. D. I. POSTON, et. al., "Design of the KRUSTY Reactor" NETS-2018, ANS (2018).
3. D. I. POSTON, et. al., "Results of the KRUSTY Warm Critical Experiments" NETS-2019, ANS (2019).
4. D. I. POSTON, et. al. "FRINK—A Code to Evaluate Space Reactor Transients" AIP Proc., 880, 449(2007).
5. D. I. POSTON, et. al., "Predicted Performance of the KRUSTY Reactor" NETS-2018, ANS (2018).