NUCLEAR PROPULSION SYSTEMS ENGINEERING*

W. W. MADSEN, J. E. NEUMAN, and D. H. VAN HAAFTEN
IDAHO NATIONAL ENGINEERING LABORATORY
IDAHO FALLS, IDAHO
(208) 526-0024

ABSTRACT

The Nuclear Energy for Rocket Vehicle Application (NERVA) program of the 1960's and early 1970's was dramatically successful, with no major failures during the entire testing program. This success was due in large part to the successful development of a systems engineering process. Systems engineering, properly implemented, involves all aspects of the system design and operation, and leads to optimization of the entire system: cost, schedule, performance, safety, reliability, function, requirements, etc. The process must be incorporated from the very first and continued to project completion. This paper will discuss major aspects of the NERVA systems engineering effort, and consider the implications for current nuclear propulsion efforts.

INTRODUCTION

Rocket propulsion in the 1960's was mostly a build-them-and-break-them proposition. Rocket motors were indeed analyzed, but ultimately, these motors were built and tested. Failures were noted and fixed and, after a number of successful tests, a mission-rated or man-rated system emerged. This required a large number of tests but was often the most prudent method at the time, considering all the constraints. However, we realized early in the nuclear propulsion program that the basic build/break mode was neither practical nor desirable. The cost of the propulsion system was very high compared to any chemical system; more importantly, the results of a failure of even one of these engines could be disastrous. Thus, a better method to obtain a flight-ready, man-rated system was needed. This paper will give a historical overview of that process.

The Nuclear Engine for Rocket Vehicle Application (NERVA) was a nuclear rocket to be used primarily for manned and unmanned missions to the moon and Mars. A number of engines were tested, most at a power level of 1100 to 1500 MW; the largest engine tested produced 4000 MW. The NERVA engine was nearly fully-developed, resulting in basically flight-ready hardware that had been tested at full power for over one hour with altitude simulation of 70,000 to 100,000 feet. These tests included the complete engine system, with controllers, instrumentation, prime movers, drivers, and valves, or in essence, the entire flight engine. The complete engine system was started, restarted, shut-down, pulse cooled and, in general, exercised to the full limits of operation.

*Work supported by the Exploratory Research and Development Program of the U.S. Department of Energy, Idaho field Office, under DOE Contract No. DE-AC07-76ID01570.

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

SYSTEMS ENGINEERING AND ADHERENCE TO REQUIREMENTS

The entire NERVA project was driven by a top-down approach with the requirements driving the system. Using the engine requirements as guides, sensitivity studies and operational maps were computed for all components and subsystems including controllers and test facilities. Thus, each component was given operational constraints such as delivered flow, operational neutron flux, temperature, pressure, etc. Each constraint had an associated tolerance and confidence level. Each component was assigned risk and reliability numbers. It was important that both the system and component personnel knew, from the sensitivity studies, how important each component was to the overall system requirements from all aspects. We could thus determine the effort (money, time, testing) to be expended on each subsystem, and the effect on the system if any component failed to meet its requirements. More importantly, we could see necessary trade-offs or component modifications. For example, we determined early that the turbopump was only critical at start-up and shut-down. Full power operation was relatively easy for existing high performance turbopumps. Thus, instead of developing all-new hardware, we modified an existing Titan turbopump by removing the last turbine stage and changing the pump inducer. On the other hand, nozzle development was very critical, since we had to operate a nozzle with pure hydrogen on both sides of the tubes. This had never been done before (or since) at these temperatures. During start-up, the nozzle chamber chokes as expected, but two-phase hydrogen also chokes in the nozzle tubes, actually going slightly supersonic. This loads the pump and pushes it toward stall. Because the systems engineering used a top-down approach, this was anticipated to be a problem area, so appropriate instrumentation and tests were devised to accurately evaluate these effects.

Analytical Tools and Testing

As stated earlier, we realized that cost and practicality would not allow the business-as-usual approach. Therefore, we set out to integrate all the components with the entire system so that optimum use could be made of testing and analysis to achieve the goals of the program with the least number of full system tests and with minimum risk. This involved the development of analysis tools. The verification of these tools had to be more than the benchmarking commonly done for computer models. All NERVA models were verified by both component and system tests. Analytical tools were always benchmarked anytime there were computer software, firmware or hardware changes. The combination of analytical tools and tests was used to determine component sensitivities and operational maps. When system-level design parameters imposed constraints on the components, the operational maps were adjusted accordingly. Trade-off studies were carried through many iterations. From this, a system emerged in which, if each component operated within its operational constraints, the entire system would do the same. Finally, as part of the required verification, the analytical models had to predict test behavior before the test was run to be considered a data point for statistical calculations.

Very close cooperation was required between the component and system testing staffs and the systems engineers developing the analytical tools. Each component test had system analysis instrumentation included in the test, and was designed to test the component's interaction with the system, as well as basic component performance. This greatly accelerated the development of analytical tools, since verification could proceed before all-up system tests were conducted. There was

also synergism with the component developers since the required instrumentation type, accuracy, sensitivity, etc. could be determined before each test and the component operation map determined with a minimum number of tests. This synergism led to many success stories within the NERVA testing program and several of these are included below.

The turbopump is critical in its off-design performance. A series of tests was run to determine operation using liquid hydrogen with low or negative Net Positive Suction Head (NPSH) pressures. The turbopump designers were not particularly interested in this, but the system designers knew that it was vital, since a critical operational mode of the system would involve just these conditions.

Systems engineers used the nozzle tests for determination of specific impulse, Isp. Load cells, flowrate measurements (at least 4), and accurate chamber pressure and temperature measurements were added to all nozzle tests. Again, the nozzle designers were more concerned in keeping the nozzle together, but the instrumentation was required from a system standpoint to determine engine Isp. The instrumentation required in this case included very expensive tungsten/tungsten-rhenium thermocouples which were necessary to the required accuracy. The data from these tests have been valuable for estimating Isp of concepts being considered today.

The nuclear system must be able to bootstrap, and, in the case of the NERVA engine, the main source of energy for the bootstrap comes from the nuclear reflector, so a series of tests with a reflector assembly was run using over 500 thermocouples in the reflector. This was not required by the reflector designers, but was essential for the system designers.

These are just a few examples of how synergism was established between component and system tests and how the type, number and accuracy of all instrumentation was determined by systems requirements. Since instrumentation costs varied with the level of sophistication required, this minimized costs for instrumentation but still ensured that the chosen instrumentation met the requirements for the testing. Test operation and methodology were determined analytically to make sure the maximum amount of information could be obtained during each test. A number of tests indicated unexpected system characteristics, and special tests were conducted to fill in the missing information. Most importantly, comparison of pre-test predictions with actual data was most helpful in determining analytical deficiencies. This process was greatly enhanced by having the same people involved with the analytical tool development, pre-test predictions, test operations, "quick-look teams", data analysis and model development. This was essential to the integration of all components into a coherent system. The products of this testing effort were analytical tools predicting such properties as temperature, pressure, flowrate, reactivity, and heat flux. These predictions were known to be correct within a known tolerance band and to have a specified confidence level, e.g., a tolerance of +2% or -3%, 95 percent of the time. Only when the tolerances and confidences were known was a code considered verified.

The advantages of such analytical tools became apparent when they were used successfully to predict such occurrences as turbopump oscillations, impeller rubbing, nozzle tube holes, control valve failure, and instrumentation failures during the tests. When test conditions did not match what was planned, we also did computer runs using the actual test data for input. Instrument-to-instrument predictions were compared by the computer to show variations that might not be spotted otherwise. Such interplay between testing and analytical tools led to great confidence in the components and system as a whole.

Analytical Tools and Configuration Controls

Analytical tools consisted of simulations of all components individually as well as of the entire system. For all large analysis efforts, such as the complete engine, we had at least three different tools: a very comprehensive digital model, a simplified digital model and an analog (later hybrid) computer model. These were used together and the results checked against each other. Usually, the simplified models were for trend predictions; the hybrid and simplified digital models for controller analysis; and the large complex digital model to confirm and refine results. Analysis of model sensitivities to such things as heat transfer coefficients or node arrangements was conducted to determine the level of detail required in the model. The codes were not made more complex simply as a matter of course. Often the big code is not the correct one for the job. For example, a two-dimensional neutronics model is only needed if there are major precursors in the second dimension. All codes have certain tolerances but these are rarely published, even if known. Code results are, too often, believed without question and without regard to accuracy or confidence level. More detail does not necessarily make a more accurate model.

Each computer code, including each input deck or model, was under strict configuration control. A user could change any part of the model or code, but these changes were printed out at the top of each run. These runs could not be used for official output, but could only be used to help guide the code development effort. Once a proposed modification was verified and documented it was brought to the change board (a panel of cognizant and responsible professionals) for incorporation into the system if the board agreed. Computer codes and analytical models were write-protected; they could be read and used but not changed except by the change board. We thus had standardized analytical tools as well as models. Analysis run in previous years could thus be repeated and checked, a feature that was used frequently.

Properties used for analysis were also controlled by committee for all materials and fluids used. Thus, all analysis was made consistent by using this standard set of properties. A report was published and updated once or twice a year as new or updated data were obtained. The importance of this can be seen by realizing that when we started we did not know the spin state (para- or ortho-) of the hydrogen in the core. We knew that in the cryogenic conditions of the tank we had para-hydrogen and, if equilibrium conditions were maintained, we would have ortho-hydrogen at the core exit. We ran special tests to determine the hydrogen properties. As it turns out, the hydrogen transit time through the core is faster than the spin state transfer time. Therefore, we had para-hydrogen throughout the system even at a core exit temperature of 2600 K.

NERVA ENGINE OPERATION

Figure 1 is a typical operation map for the NERVA engine. The x-axis is chamber pressure and the y-axis is chamber temperature. The map shows lines of constant engine flowrate (lines down left to right), constant reactor power (lines upward left to right), and the engine operational limits. The top limit line is determined by the temperature at a point about 2/3 of the way down the core and 4 to 5 inches in from the core outside radius. This is the hottest spot in the core at full power because the neutron flux peaks at this point due to the reflectors and the control drum location. The right-hand constraint is determined by the turbine power control valve wide open condition, resulting in maximum power to the turbomachinery.

Figure 1. NERVA Performance Map

It should be noted that the core is not energy limited, but is limited by power density. More power can be generated by just enlarging the core; we did this for the Phoebus design in which the core was increased in size to obtain approximately 4000 MW rather than the NERVA nominal 1100 to 1500 MW. Power can also be increased by modifying the fuel for increased power density and leaving core size fixed, but material temperature limits may be encountered. Reduction in core size is limited by the requirement to maintain core criticality.

On Figure 1, the left-hand limit of the operational map is determined by pump stall, followed closely by structural limits. As engine operation goes to the left, flowrate is reduced, but since temperature is held constant, pump head rise requirements do not lower at the same rate. Thus, the pump will eventually stall by induced cavitation on the impeller. For concepts other than NERVA, this limitation can be lessened by reducing operational chamber pressure resulting in reduced pump operational requirements, or removed entirely by going to a pressure feed system. There is still, however, a structural limit on that side.

The lower left quadrant of Figure 1 has been extracted to show a typical NERVA start-up in Figure 2. Beginning at the lower left corner, the engine is first conditioned by running hydrogen through the system to chill the pump. Next, at 0 sec, the bootstrap is started by ramping up both the chamber temperature and chamber pressure demand. At 3 sec, flow is introduced into the turbine, and the reactor power is increased. This is the most critical part of the start-up procedure. As can be seen, start-up at this point does not follow the straight line (indicated by dashes) because of the necessity of avoiding stall and various material limitations on the left, and reactivity feedback problems on the right. At approximately 20 percent power (11 sec) the engine is low in temperature and high in pressure. The start-up line then curves to the left, with chamber pressure increasing slowly

Figure 2. NERVA System Start-up Procedure

while chamber temperature increases more dramatically. This avoids reactivity feedback problems which can occur especially if a restart is required. From 2200 °R and 259 psia, the track to full power then is basically straight with the controllers set to avoid an overtemperature condition as full power is reached.

Shut-down is initiated by reducing chamber pressure demand while holding chamber temperature at full operating conditions. This reduces reactor power while still achieving, or even increasing, engine Isp. After a small hold at the reduced power (usually 20 to 30 percent), chamber pressure and temperature are reduced following the pump stall line until the reactor is scrammed. Hydrogen flow must be continued to keep the engine from overheating due to reactor decay heat. The lower power scram is required to reduce the total integrated hydrogen required during this cooldown phase, since decay heat is most directly related to reactor power at the time of scram. Hydrogen used during cooldown is producing thrust, but as can be seen, the chamber temperature for engine conditioning, start-up, and shut-down after scram is considerably lower than that at full power or during power retreat. Thus the engine delivers lower Isp at

this time resulting in lower average Isp for the entire operation than that obtained at full power. This has to be taken into account for any engine concept since it effects overall thrust and Isp. For nuclear propulsion, full power Isp of, for example, 1000 s will result in an overall burn Isp of about 980 s, depending on burn time and scram power.

In analyzing and comparing engine concepts proposed since the NERVA program, all operations must be optimized for safe and reliable operation with maximum overall performance. The start-up and shut-down are very critical to engine operation. Sound systems engineering analysis must be used now, before down-selection of a concept, to determine these constraints. Waiting to do so on an operating reactor will most likely result in disaster. The analytical tools used must be verified to make sure computed sensitivities truly represent real hardware operation.

LESSONS LEARNED AND FUTURE APPLICATIONS

NERVA engine operation was streamlined by use of the analytical tools developed during the program. On the Hot Bleed Engine, a hot bleed port control valve was eliminated by changing the injection angle of the holes in the hot bleed port, converting it into a fluidic valve. Low- or zero-NPSH start-ups were achieved by softening the bootstrap of the engine. Turbine power bleed was changed from the hot bleed port to a dual topping cycle, in which hydrogen for the turbine drive was taken from the peripheral shield exit, while turbine exhaust was returned to the reactor through the top shield. At first it was thought that there was not enough energy for this arrangement, but there is actually too much energy if turbine bypass control is used. The computer models predicted that if too much cold hydrogen was introduced into the core, the control drums would be ineffective. As a result, we were prepared for this when, during one series of X engine testing, we were running a shut-down and too much cold hydrogen inventory was allowed into the core. The result was an engine operating at 250 MW with the drums fully inserted. The hydrogen was shut off momentarily, reducing the cold hydrogen inventory, and the shut-down proceeded without any problems. The same type of problem was evident during start-ups. After two abortive starts on the X engine, a large quantity of cold hydrogen was in the core during the third attempt. The drums went to their fully closed position but instead of scramming, the reactor continued on a 0.1 second period start ramp. The drums came off the pegs at approximately 300 MW, control was regained, and the start-up proceeded.

Although the details of power levels and control difficulties vary with engine design, all nuclear propulsion concepts have the same generic problems with start-up and shut-down. Full power operation is far easier than either start-up or shut-down, both of which have relatively narrow windows of operation. These windows are also concept-dependent and for NERVA the narrowest portion of the window was near 30 percent power. The windows will become increasingly narrow if these fully-enriched cores are made smaller and the power density is increased. At full power, the hydrogen is usually hot enough to act as a slight poison for most concepts, but at reduced power this is emphatically not the case. All these things must be considered to make a viable operational nuclear propulsion system.

In the NERVA engine, the tie rods holding the core together were surrounded by hot fuel elements. The tie rods were cooled with flow taken from above the core support plate and dumped back into the system below the core support after passing through the tie rods. On the NRX A-3 full-power burn, one of the

turbopump rpm indicators erroneously indicated overspeed conditions due to a sensor problem. This automatically scrambled the reactor, and the sudden cessation of flow caused the tie rods to overheat from an operational temperature of 500 K past their upper limit of 612 K and off the top of the recorders at over 1000 K. Fortunately, a leaky check valve located just below the pump discharge had stored a sufficient quantity of liquid hydrogen at high pressure to cool the system and prevent failure of the tie rods and nozzle tubes. The core was restarted after determining that there were no holes in the nozzle tubes, and we were near 50 percent power before the control drum rollout was excessive. We ensured that this capacity would always be there by drilling holes in the check valve. This type of passive emergency flow system is useful for all nuclear thermal propulsion system concepts. The NERVA core was enriched to better than 90%, and the energy available was considerably higher than that available in the incident at Three Mile Island. Had this been a water-cooled reactor there would have certainly been a non-nuclear explosion. There was material degradation but no serious loss of fuel with the flow stoppage test. We did lose part of the core on restart. Though we learned a lot about system safety, it will be imperative that such occurrences not happen in a new nuclear propulsion test program.

Changes or improvements in components or operation can readily be determined when the systems engineering is being performed correctly. An example of this is that the tie rod temperature problem resulting from a reactor scram at full power was eliminated by taking the tie rod flow from the pump discharge. This gives the passive emergency flow system quick access to a critical part of the core. Another example was the hot bleed port turbine drive system, which was costing Isp performance since it put some of the hydrogen overboard at turbine discharge temperature of 300 K instead of at the chamber temperature of 2600 K. This loss was eliminated by taking the turbine drive flow from above the internal shield and returning the turbine exhaust below the core support plate. The hydrogen stream was heated to full chamber temperature as it passed through the core. In a further attempt to make a simpler, more reliable system, we are now considering concepts which require little or no turbomachinery.

One NERVA test, conducted with very little previous analysis, was appropriately named the TNT test. The test was to determine the consequences, due to the highly enriched nature of the core, of a full reactivity insertion while operating at full power. It was felt that no amount of analysis would satisfactorily answer this question, so it was decided to run an all-up test. The TNT test did answer the question with a minimum cost, since an old reactor was available, and produced unquestionable results (an advantage over analytical modeling alone.) This test was conducted by running a rail spur line out in the desert to transport the reactor. Explosive drives were put on the drums to operate them a conservative 1000 times faster than would occur during any mechanical drive failure. The core was brought to full power and the drives actuated. A relatively energetic explosion occurred and was photographed. (This is easily the most misunderstood event of the program. Many erroneously believe to this day that it was nuclear in nature, while it was caused by a rapid gas expansion due to the large stagnation temperature increase in the core.) Though the test was drastic, the question was answered with testing when it could not have been with analysis alone.

Sensitivities are usually computed by exercising an analytical tool and changing a single variable while holding all the other inputs fixed and noting the effect on system performance parameters. This is a valid but time-consuming effort resulting in very large matrices, especially if second- and third-order effects are included. This task can be reduced considerably by use of second law

analysis developed for NERVA. This technique can determine the slopes of all the sensitivity curves at any point by correct comparison of irreversibilities, optimization factors, and availability functions for all system and component drivers. This method takes only a few computer runs and greatly reduces the decision matrix size by quickly eliminating those parameters with very flat slopes that thus have little effect on system. In all cases, those primary variables indicated by the second law analysis were confirmed by using the standard perturbation techniques.

Systems Engineering and a New Nuclear Propulsion Program

Future analysis should step back another level from what we used on NERVA. The overall requirements for the propulsion system should be determined first. NERVA systems analysis wasn't allowed to be comprehensive enough because the mission requirements had been set independently. At the INEL, we have run mission maps for all the planets and looked at mission profiles.^{1,2} These studies allow us to rank the main engine requirements derived from mission requirements as follows: engine Isp, tank mass fraction, operational characteristics (clustering, staging, use of perigee pulsing, degraded engine operation), total thrust, and engine thrust-to-weight ratio. Though Isp is the top priority, there are limits to the importance of that parameter. Figure 3 shows that for a manned Mars mission, an Isp larger than approximately 1200 s does not reduce the mass ratio in low earth orbit significantly. It does broaden the launch window, but this has to be weighed against risk, cost, schedule, or reliability impacts resulting from this increased performance.

Figure 3. Specific impulse sensitivities for Earth to Mars

Figure 4 shows the effect of engine thrust-to-weight ratio. Our analysis of mission profiles indicates that optimum engine thrust is approximately 15000 lbf. This size engine will capture all profiles in the NASA Mission Profile⁴, either as a single engine or clustered, as for a manned Mars mission and its precursors. Preliminary reliability analysis indicates that optimum cluster arrangement is three to five engines. This arrangement also makes the earth departure burn safer and reduces or eliminates gravity losses by allowing optimum perigee kicks. We have developed computer codes to optimize the perigee kicks taking into account masses, gravity losses, and burn duration and timing. Radiation, primarily from the Van Allen belts, is also included to take into account overall mission safety.

Figure 4. Thrust and thrust/engine weight sensitivities for IMLEO = 500 tonnes, $I_{sp} = 1000$ sec, and $C_3 = 16 \text{ km}^2/\text{s}^2$

Concept selection and development will not only depend on the mission requirements for the system, however. Because our world is more environmentally conscious today than it was during the NERVA program, the testing requirements will also drive the engine design. A first-level study of the testing requirements has been done at the INEL.³ It concluded that the requirement that the engine be tested forces any concept to conform to the realities of physical practicality and economic feasibility. In general, lower thrust engines (less than around 20,000 lbf) will be much easier to test than high thrust engines (higher than about 70,000 lbf.)

Analysis and testing are of equal importance to a nuclear propulsion program. They must both work in a coordinated effort to minimize cost and schedule, and if done correctly, the end product will be a flight-ready engine with verified analytical tools to back up decisions concerning everything from component selection to engine operation. All the "why" questions should be answered: Why this turbopump? Why this nozzle? Why operate in this way? Also,

the operational limits on all components and the overall engine operation will be understood. Basic research will continue as engines are developed; as technology improves, the impact of these improvements can readily be assessed and decisions can be made quickly as to the practicality of including them into the on-going engine design.

The NERVA program resulted in flight-ready hardware which was obtained with a minimum number of tests. All engine concepts have problems and strengths, but these can only be determined by a good systems analysis effort focusing on the engine requirements and not on any particular concept. Future development must be driven by engine requirements derived from both mission and testing requirements. This is the top-down approach that leads to informed decisions which can be backed up analytically.

REFERENCES

1. W. W. MADSEN, J. E. NEUMAN, T. S. OLSON, and A. S. SIAHPUSH, "Mission Maps for Use in the Choice of Specific Impulse for Manned Mars Missions," AAS Paper 91-503, AAS/AIAA Astrodynamics Conference, Durango, CO, August 1991.
2. W. W. MADSEN, T. S. OLSON, and A. S. SIAHPUSH, "Optimization of Perigee Burns for Manned Interplanetary Missions," AIAA Paper 91-3402, Conference on Advanced SEI Technologies, Cleveland, OH, September 1991.
3. W. W. MADSEN, T. S. OLSON, and D. H. VAN HAAFTEN, "Nuclear Rocket Engines and Test Facility Requirements," AIAA Paper 91-3415, Conference on Advanced SEI Technologies, Cleveland, OH, September 1991.
4. NASA Mission Profile, Rev. 8

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

END

DATE
FILMED

7 / 14 / 93