N92-11092

NERVA UPGRADE: NON-NUCLEAR COMPONENTS

Stanley Gunn
Rockwell International, Rocketdyne Division

As Stan Borowski pointed out, the technology that did exist back in the 1960's and at the start of the 1970's under the ROVER/NERVA program was rather substantial, but there have been advances that have occurred since the initial design of the NERVA. Some of those advances were accomplished under the Phoebus program, the technology program that Los Alamos and Rocketdyne were involved with in Nevada.

Other advances have occurred in the development of the shuttle engine and related chemical rocket engines. What I would like to talk about is what would be realized if we were designing an engine today based upon the original accomplishments of the ROVER/NERVA program, but feeding in these advanced technologies; what would its characteristics be, what would it be able to accomplish?

Now to start off, I have set down some hypothetical requirements for a typical manned Mars mission. I'll try to highlight the areas that would influence the selection of the design details of the engine. As shown in Figure 1, I am assuming; 100,000 pound thrust engine with performance requirements in excess of 900 seconds; a maximum weight of 14,000 pounds without the shield, which is going to be a bit of a challenge to achieve; a full performance operating range of 50 percent thrust at full Isp up to 110 percent. Then, reflecting the concern to have a very reliable system, we had dual turbopumps with a pump-out capability that would give us a capability of operating at 70 percent of rated thrust at full Isp. (Incidentally, we ran Phoebus 2A with dual turbopumps).

Further I am going to assume that we are going to be able to engineer the pumping system so it will be able to take hydrogen as a saturated liquid from the tank, accelerate it to one velocity head (which means that we are going to be ingesting vapor), and pump it to the full requirements of the reactor in terms of pump outlet pressure and flow rate.

As far as the maximum operating time of two hours, that comes, in part, from a belief that by getting the engine thrust up to about 100,000 pounds for the typical Delta V's that we have been talking about here (fairly fast trip times), we will be able to limit the burn time, of the engine that runs the longest, to two hours. I have done that because I wanted to tie it back to what was accomplished in Nevada and the nuclear furnace, where fuel elements were run for approximately two hours.

I have assumed 6 restarts, and that takes into account using one engine for a number of maneuvers, and a transition from flow initiation to full thrust of 30 seconds. That goes with a ramp rate of something like 100-150 degrees per second coming up in temperature and thrust, with transition from 50% thrust to cut-off of about 30 seconds, and a maximum core temperature that we have to remove afterheat of 1800 degrees R.

There were a number of reactors that were examined back in the 1960's and a number that are now currently under examination. Back in the 1960's, we had the solid core reactor, which is typified by the ROVER/NERVA program, and there were some fast metallic systems that were looked at.

We did some engineering design studies of engine systems incorporating the GE 710 reactor, and also the Argonne National Lab had a similar fast metallic concept. We also did a design study with Frank Rom here on an engine based upon the utilization of tungsten 184, and the use of water moderation to provide very attractive engine cycles.

The presentation that I am limiting myself to today is the solid core reactór - ROVER/NERVA. I would like to say a couple things about the expected performance as shown in Figure 2.

We are talking now about temperatures in the range of 4500 degrees R to about 5580. The epsilons (nozzle expansion ratio) show what we can expect for the performance of a high pressure system. This includes nozzle losses from the gas kinetics and nozzle boundary layer. Divergence effects are also included. For a condition of an epsilon of 500 and a gas temperature of 4860 R or 2700 degrees Kelvin, the Isp is on the order of 920-925 seconds. You can see from this chart that there is not a lot to be gained by going to expansion ratios higher than 500. We are collapsing down to 800 to 1000 at almost the same value.

Now, if you want to look a little bit closer at how those numbers came about, there is a series of comparisons that might be of interest. A 250 K engine at a chamber pressure of 1000 psi, has a theoretical performance of 1029 seconds. If you take into account the kinetics of what's occurring in the expansion process, it drops it down to 1028 seconds. The boundary layer losses drop you down to 1014, and the divergence effects drop you down to 1011 seconds.

It's important to look at the boundary layer effects, because if we go next to a 75 K engine, where we have less flow and therefore, more boundary layer effect, we've dropped it down to 1010, with divergence of about 1007. This is all for the case of 3100 degrees Kelvin that we have examined here.

Now if we go clear on down to a very low pressure to take advantage of the increase due to dissociation and reassociation, here is what your numbers come down to. The theoretical performance is very high. But as you examine what happens in the kinetics (the recombination, relaxation), you find that you drop down to about 1372 seconds -- this one is for 7,000 degrees R, -- and you suffer losses down to about 1300 for the 7,000 case. When you take a look at that same effect based upon 3100 K, you find that the performance is somewhat higher than the 1020, but not a lot.

Now, I would like to talk a little bit about the selection of cycle. In this particular presentation, we have made Isp our "God." We are trying to find out what we can do to get the maximum Isp from the temperature, and so you will see this emphasized in the charts that follow.

One of the first things you would like to do, given a certain Isp or, rather, temperature, is to use an expander cycle. The designer has several approaches available to be able to accomplish the circuitry of the flow to get the temperature of the hydrogen, up high enough to be able to drive the turbine. The obvious reason for this is you don't want to pay a penalty in terms of the Isp by having less than full temperature in all the gas.

In Configuration A in Figure 3, we have taken a portion of flow down through the tie tubes on up to the point where it is going to join some flow that has come down through the nozzle. It joins the flow that has been split off and that goes up through the reflector.

The goal is to get the temperature of all the gas coming into the core as high as possible, to maximize the amount of heat that the fuel elements can give to the hydrogen, which will allow you to go to as high a thrust as possible. That is a key point in being able to raise your thrust-to-weight ratio: get the temperature up so that the full power of the reactor can heat more working fluid.

Now in Configuration B in Figure 3, we have done something a little differently. What we have done here is assume that we can get all the heat we need to drive the turbine through the tie tubes. This allows us to minimize the heat pickup up through the nozzle and then up through the reflector. However, remember I said I want to make Isp my "God" here. Any heat that is transferred to the nozzle up to this point is a loss. It's taking enthalpy out of the expanding gas, and it drops your Isp a little bit. So what you would like to do is make the nozzle all adiabatic, but we can't do that because of the materials.

This study is based upon a ROVER/NERVA core that makes use of a number of clusters, as shown in Figure 4. In this case, there are a total of 6 of these 19-hole fuel elements residing around a center element, resting on a core support block. The fuel element is approximately 52 inches long. The tie tube assembly is used to get the enthalpy to drive the turbine. This particular design has pneumatically driven actuators.

Now, let's take a look at what happens to a high expansion ratio nozzle if we try to design it to make use of the maximum amount of enthalpy. In Figure 5 we have assumed carbon/carbon composite as the material, and we have plotted the maximum wall temperatures, both inside and outside, as a function of area ratio .

For the design you will see in a minute, we have chosen to limit the expansion ratio at which we attached this adiabatic cooled nozzle to about 150; the wall temperature goes to about 2600 degrees R on the composite. We can be tempted perhaps to go to a smaller cooled expansion ratio, which would mean that we would be extracting less heat. We could help ourselves with Isp, but in doing so we are going to run into a problem in what we think the maximum temperature is that the uncooled noozle can handle.

Figure 6 shows the plot of some of the calculations that have been made on the heat load on this kind of a nozzle. In this particular case, we have assumed that the hydrogen comes in and flows two ways at an epsilon of about 6. A portion of the flow goes down to that 150 to 1 expansion ratio point, then back up, which gives rise to these two values of the wall temperatures shown here. Notice that the heat flux hits the maximum around the throat, and in this particular case we are talking about 40 BTU's per square inch per second. It then drops drastically down.

In the shuttle engine, we are able to withstand heat fluxes at the throat area of about 75 BTU's per square inch per second, so we could go a bit further on up in chamber pressure.

One way to get the thrust-to-weight ratio of your system up is to reduce the size of the entire assembly by going on up in chamber pressure. One of the big drivers in terms of size is the nozzle, so that if we do succeed in operating at a higher chamber pressure, that will shove up this heat flux at the throat. However, we still have some margin to deal with there.

Figure 7 shows what this thing looks like when you make the assumptions that I just talked about. Here is a nozzle assembly involving the reactor, the throat area, the point at which the hydrogen comes in and makes a pass, and a half-portion of it down through the throat. The other part goes up through the converging section where your big heat load is. Notice how big this whole assembly has gotten. I have shown it here as if it were an extendable nozzle (this is the uncooled portion), and it's translated up around the engine.

This particular size is dictated in part not only by what I was just talking about, but by the size of this interface. If you try to get a very high Isp with a conventional nozzle system (and the reason this thing looks this way is because we have tried to avoid mach lines; shock losses in that expansion process to minimize Isp again) you are talking about a very large assembly.

There may be ways to get this down to a more manageable size. One can think about the idea of the collapsible drinking cup you take on camping trips, and put several interfaces there to be able to pull down this size. There are also some other nozzle concepts that lend themselves to better packaging, but if you go conventional and you go for maximum Isp, this is what you are faced with.

I mentioned earlier that our interest was in having a reliable system so that if we should have a failure, or sense an incipient failure of one of the turbopumps, we could continue to operate and get the mission accomplished or at least retreat gracefully to some sort of a recovery plan. If you try to provide for dual turbopumps to do this, you may have a very complex system involving a number of valves that can isolate the pumping, or that can also get the pumps to share the load equally.

In Figure 8, which we have patterned after our experiences in Nevada, we had actual flow pumps that had negative HQ curves. That kind of a system is inherently self-balancing. The pumps share the load equally.

We have done it here by picking a design point that is close to the design specific speed line, but is far enough over from the predicted stall region so we can actually throttle at full temperature down to 50 percent. This is the value that I assumed in my example. It also represents the kind of limit that the reactor people are comfortable with in terms of having full temperature, but reduced coolant flow going through. If you suffered a failure of a turbopump, you would move to a new operating point out to the right, where the developed head and the flow rate intersect with the reactor load line.

This example shows an ability to meet that requirement, but notice that we are getting close to what was called a negative flow incidence. That's the case on the inducer at the front end of the pump where the flow incidence angles on the impeller, the front end of the impeller, goes negative, and then your NPSH requirements come up. So if you want to try to operate this pump at 70 percent thrust (for reasons of retrieval on your mission) and you want to operate with negative NPSH, it may be necessary to add a boost pump that would be hydraulically driven so you could match the speed to give you the proper incidence angle.

Figure 9 is a cartoon of one version of the pump that could do this. In this particular case we have patterned it after what we did on the Mark 25. The design incorporates hybrid hydrostatic bearings, at the outboard end of the turbine, which were proven in Nevada to be able to operate in a very satisfactory manner.

We wanted to go for the hydrostatic bearing to get rid of any materials that were in the bearings or anywhere else in this turbopump that would suffer any kind of damage from the intense nuclear environment that we anticipated. Those bearings actually provided the means of doing that, and they worked.

We also tested an advanced inducer which actually went up in flow capacity by 50 percent area. It was made out of titanium so that we could keep the weight overhang off the stub of the shaft reasonable, to maintain a critical speed where we wanted it. We also reduced the incidence angle, a design Q over N, to about 1.5 degrees, and we tested it. We found that we were able to ingest not only a saturated liquid going into the pump, but we kept going and we found that we could ingest up to 30 percent vapor and still the pump put out full pressure.

If that can be realized in a flight system, it enables us to pump a saturated fluid from a tank. The other thing we have shown is a single stage turbine over at the drive end. We could play the games of going to two stages there, take the pressure ratio and adjust it across each stage and make more effective utilization of the turbine drive fluid, depending upon the temperature that we are able to put into that drive fluid by the tie tube circuitry. So, we have some flexibility there.

We have also looked at integrated, pneumatic fluidics control systems to come up with a control system that would enable us to operate the entire engine in an intense radiation environment (see figure 10). Based upon the development work that was accomplished, it looks like we could do it. In other words, it would not be necessary to shield this engine from anything it does to itself radiation-wise. If you want to put a shield in this engine, it would be to protect the crew, but not because the engine requires it.

Figure 11 shows the final version of what an expander cycle engine system would look like in terms of its operating conditions. This particular setup allows us to meet the requirements I talked about, except that we are talking here about a weight of about 18,000 pounds and not 14,000 pounds.

How do we get that thrust-to-weight ratio up? Obviously most of the weight is in the reactor, but there may be a limit to what can be done there to make the weight as low as you would like to. There is another way to get that thrust-to-weight ratio up, and that is to get more thrust out of this configuration. And the thing that determines what you can get out of this engine is not the design of the pumps, not the design of the nozzle, but the power density in the fuel elements.

This particular design at 100 K has a power density of 1 megawatt per fuel element, which was actually demonstrated in the Phoebus program. There are some indications that you can get as high as 1 and a quarter megawatts per fuel element. That would raise the same engine to 125,000 pounds of thrust. That's the route that you need to examine: how hard you can push the fuel elements in power density for a given outlet temperature and a given total operating time?

TYPICAL MANNED MARS MISSION NTR PROPULSION REQUIREMENTS

• THRUST, NOMINAL	100,000 LBF
• PERFORMANCE	≥ 900 SEC
• MAXIMUM WEIGHT	≤ 14,000 LBS (WITHOUT SHIELD)
• FULL PERFORMANCE OPERATING RANGE	110% → 50%
• EMERGENCY THRUST (W ONE PUMP OUT)	70%
• NPSH (MIN)	ONE VELOCITY HEAD FROM SATURATED LIQUID IN TANK
• MAXIMUM OPERATING TIME	2 HOURS
• NUMBER OF RESTARTS	≥ 6
• TRANSITION, FLOW INITIATION TO FULL THRUST	30 SEC
• TRANSITION, 50% THRUST TO CUT-OFF	30 SEC
• MAXIMUM CORE TEMPERATURE, AFTER HEAT	1800°R

Figure 1

Performance* of NTR Operating with High Pressure Hydrogen

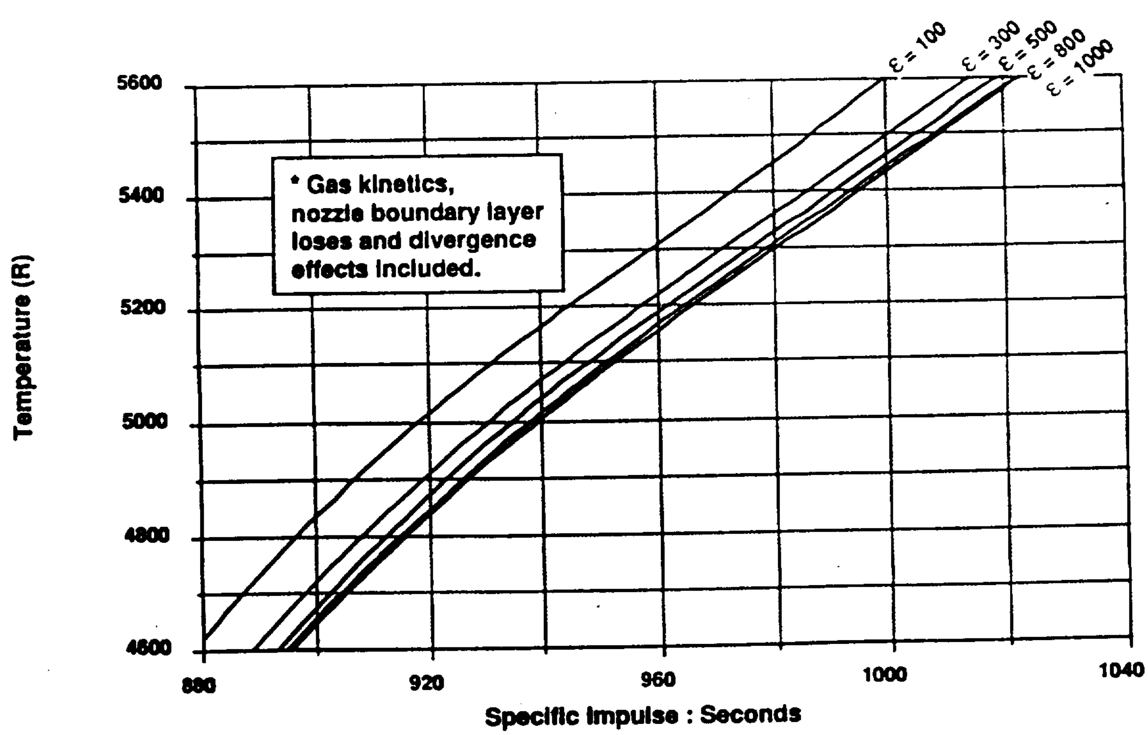

Figure 2

NTR Expander Cycles (Typical)

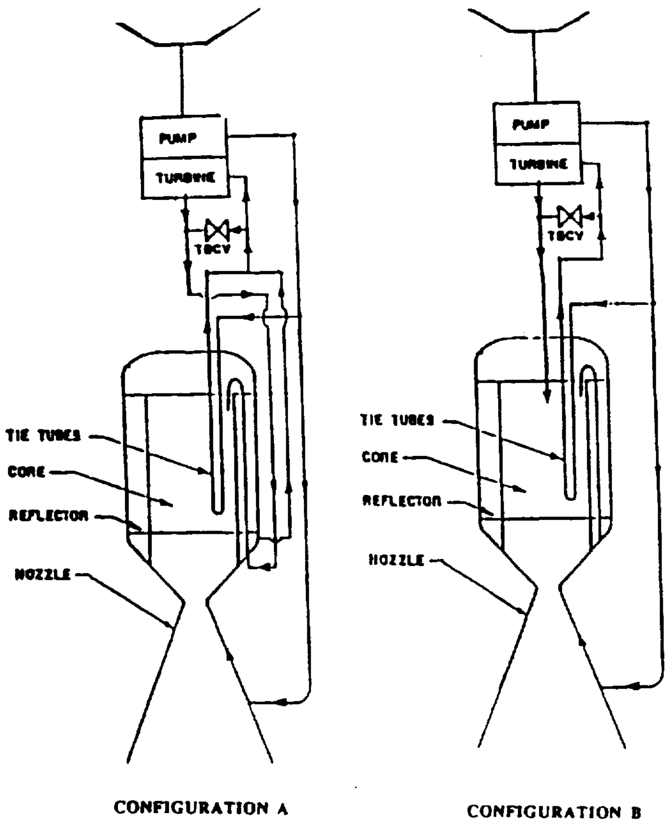

Figure 3

Rover/Nerva Core-Reactor

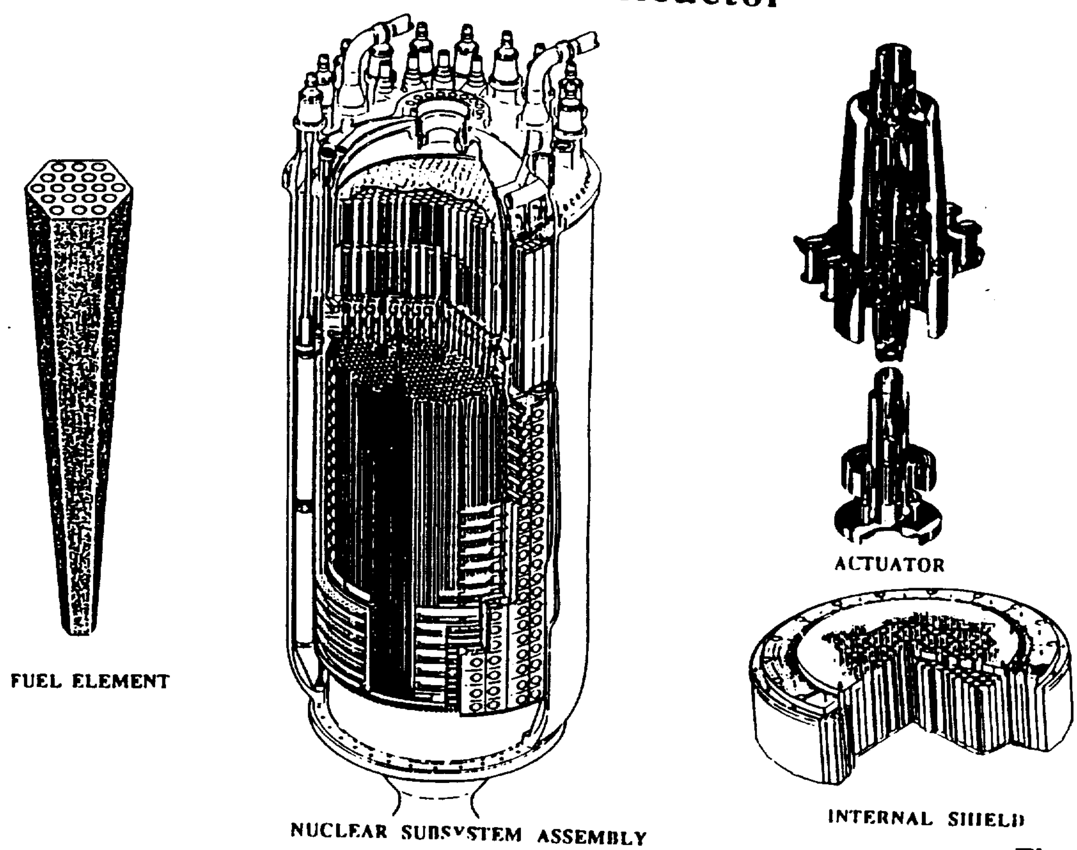

Figure 4

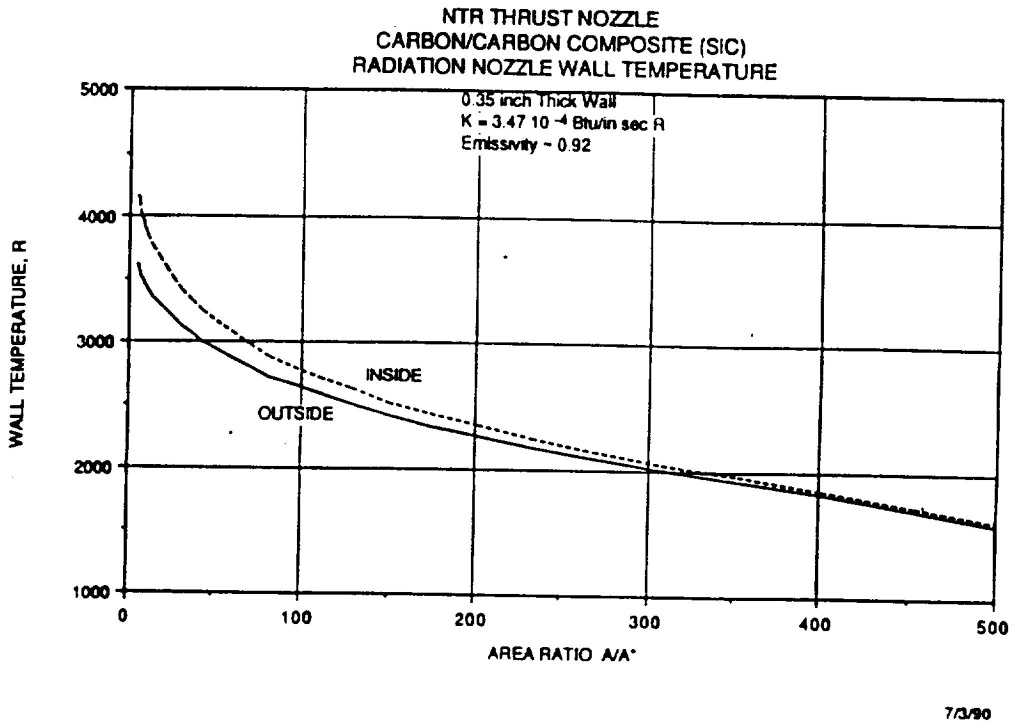

Figure 5

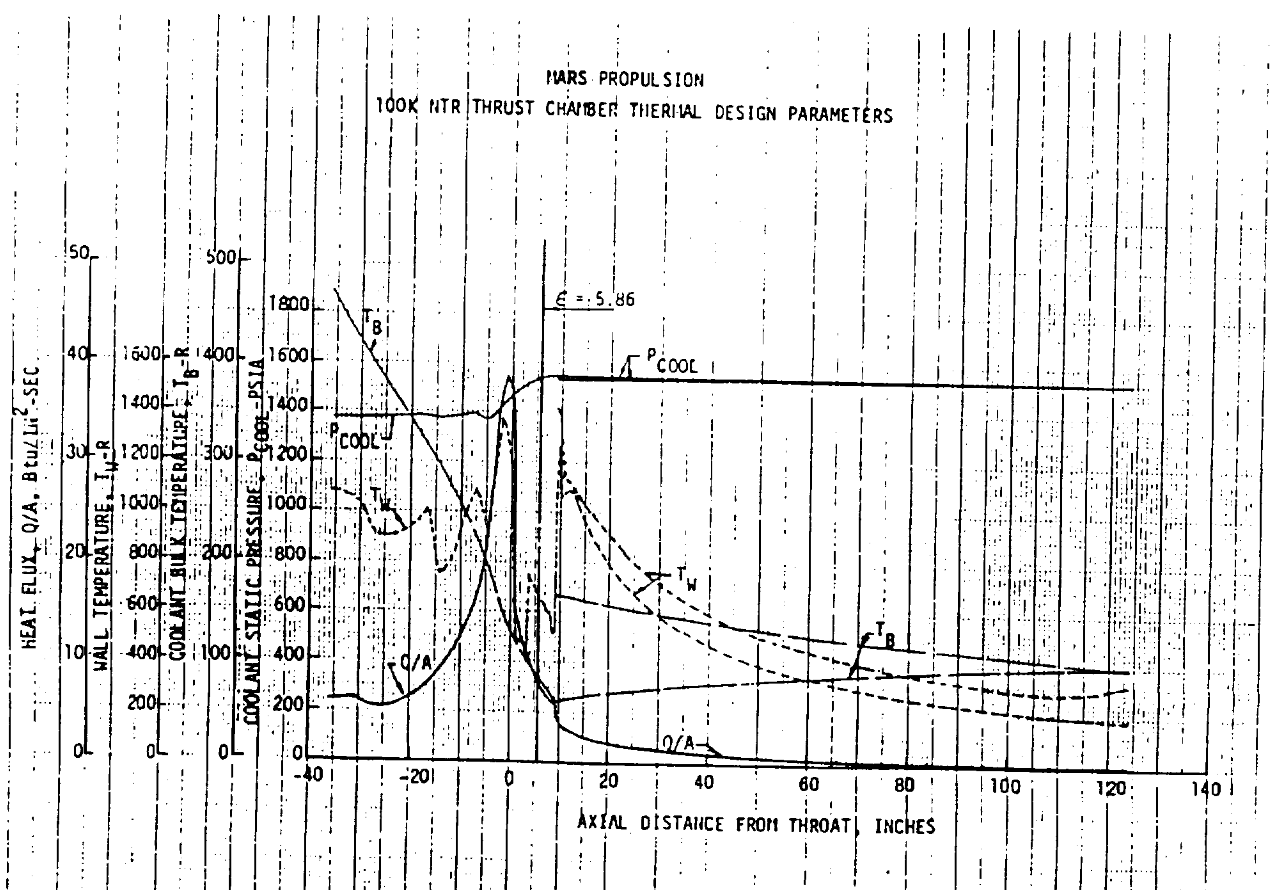

Figure 6

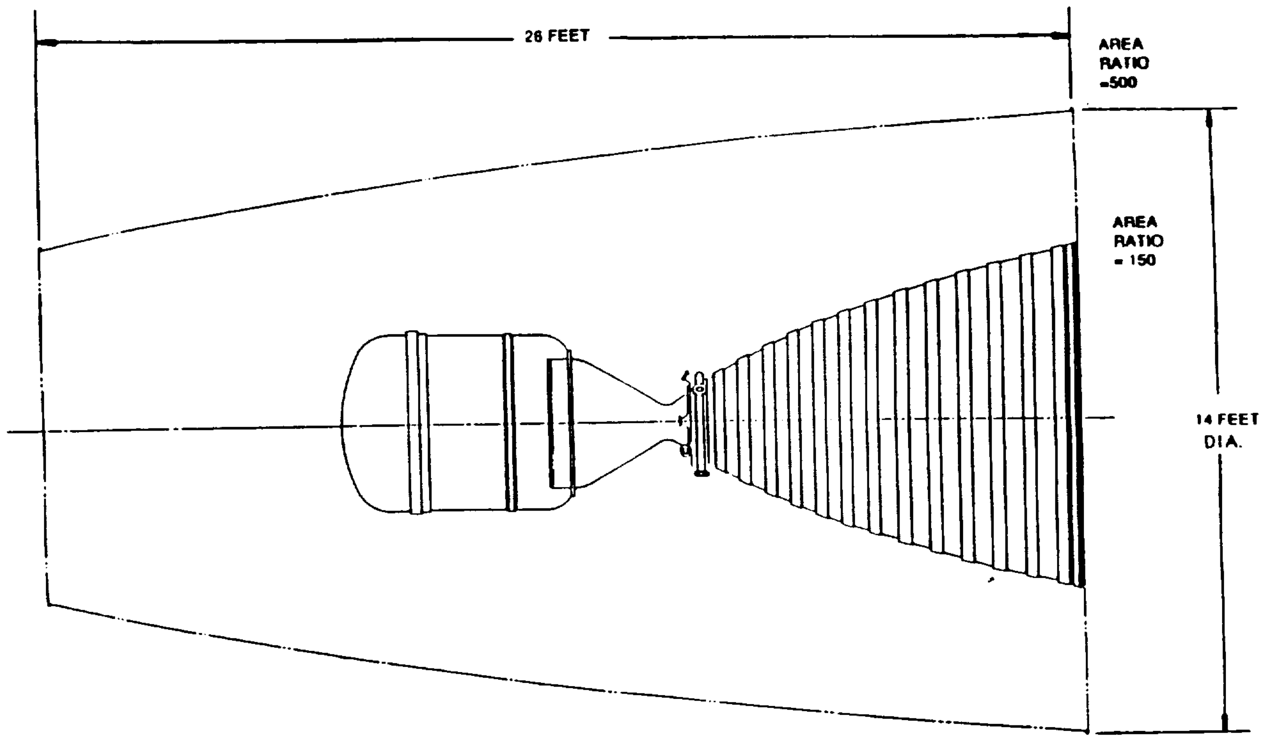

100K NTR NOZZLE ASSEMBLY

Figure 7

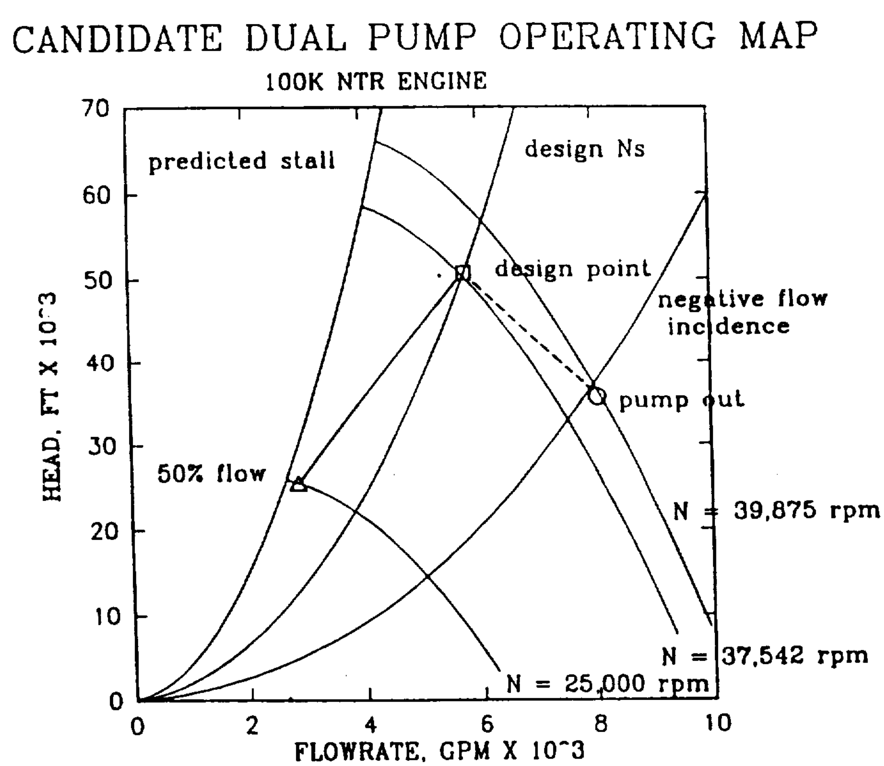

Figure 8

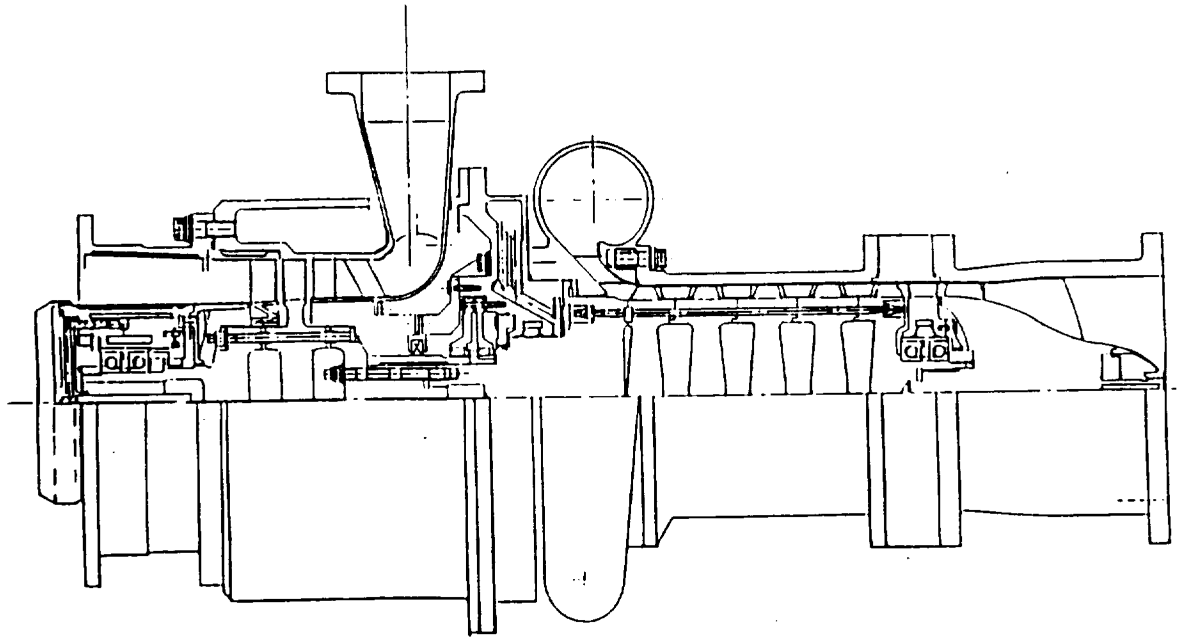

SCALED 5 STAGE MK25 PUMP
SINGLE STAGE EXPANDER TURBINE
PUMP FLOWRATE 54 LB/SEC
RPM 37,500
PUMP INLET DIAMETER 5.71 IN.

Figure 9

NTR Integrated Pneumatic-Fluidics Control System

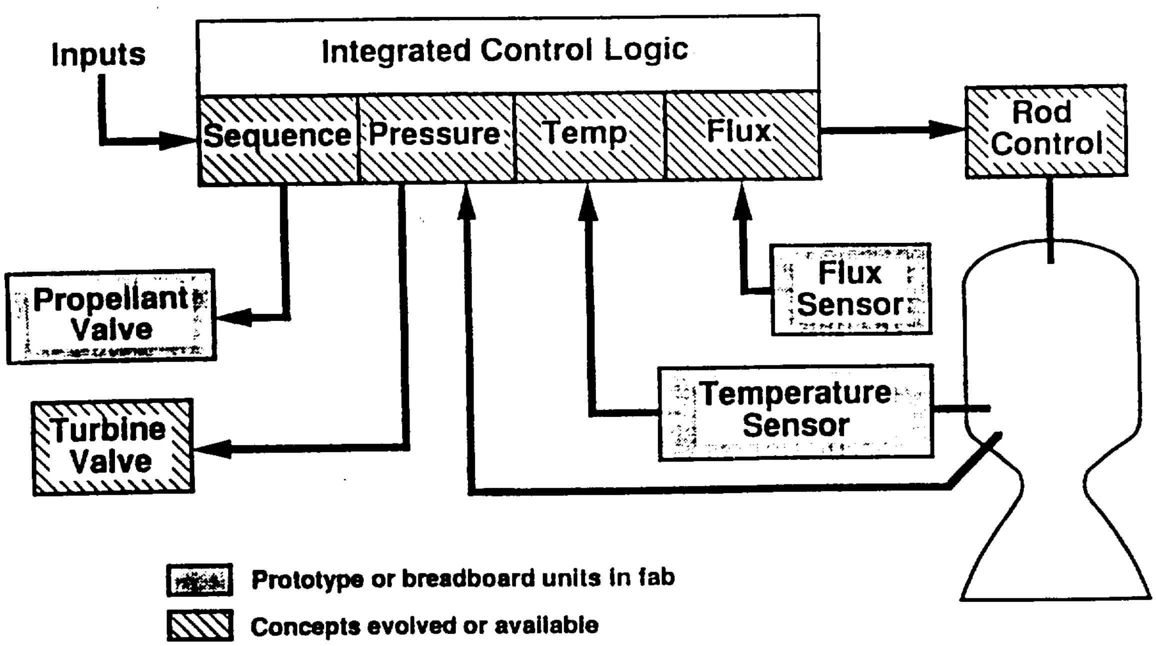

Figure 10

100K FLIGHT THRUST MODULE/FEED SYSTEM MODULE ASSEMBLY

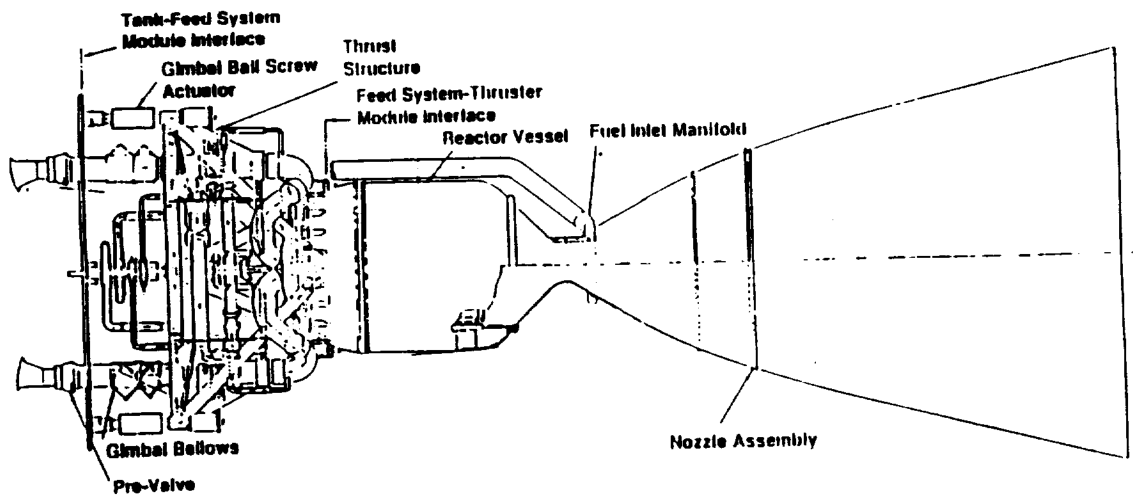

Is (ϵ =500:1, UNCOOLED SKIRT) 922 SEC
WEIGHT..14,000 LB
FLOW RATE..108 LB/SEC
ϵ (COOLED SKIRT)..................................150:1
POWER DENSITY....................................85 MW/FT^3
ENGINE CYCLE.......................................EXPANDER

Figure 11