Low-Cost Propellant Launch to Earth Orbit from a Tethered Balloon

Brian Wilcox
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena CA 91109
(818) 354-4625
Brian.H.Wilcox@jpl.nasa.gov

Abstract - Propellant will be more than 85% of the mass that needs to be lofted into Low Earth Orbit (LEO) in the planned program of Exploration of the Moon, Mars, and beyond. This paper describes a possible means for launching thousands of tons of propellant per year into LEO at a cost 15 to 30 times less than the current launch cost per kilogram. The basic idea is to mass-produce very simple, small and relatively low-performance rockets at a cost per kilogram comparable to automobiles, instead of the ~25X greater cost that is customary for current launch vehicles that are produced in small quantities and which are manufactured with performance near the limits of what is possible. These small, simple rockets can reach orbit because they are launched above ~95% of the atmosphere, where the drag losses even on a small rocket are acceptable, and because they can be launched nearly horizontally with very simple guidance based primarily on spin-stabilization. Launching above most of the atmosphere is accomplished by winching the rocket up a tether to a balloon. A fuel depot in equatorial orbit passes over the launch site on every orbit (approximately every 90 minutes). One or more rockets can be launched each time the fuel depot passes overhead, so the launch rate can be any multiple of 6000 small rockets per year, a number that is sufficient to reap the benefits of mass production.

This approach might allow NASA to have a "propellant rich" exploration architecture that makes it possible to reduce the performance requirements on all space hardware launched into LEO, since there will be abundant propellant to move that hardware out to the moon, to Mars, or beyond. This heavier, lower-performance hardware can be much less expensive and much more reusable than would be possible if so much propellant were not available. As a result, this approach might reduce the overall cost of exploration by an order of magnitude.

Think of the tethered balloon as a "flagpole" (Figure 1). The balloon supports a pulley that has the tether looped over it, and winches on the ground can lift the rocket launcher up to the stationary balloon, much as a flag is hoisted up a flagpole. A reloaded launcher goes up to the balloon from one winch as an empty launcher is lowered from the balloon to the other winch. By physically separating the two winches on the ground, the tether lines will not get tangled and the rocket can be dropped freely between the two tether lines for a few seconds before firing so that neither the balloon nor the tether are put at risk of being incinerated by the rocket.

1. Introduction

Many mission studies of alternative approaches for extending human reach into the solar system have shown that the lofting of propellant mass into Earth orbit is a dominant cost of any such effort.[1] The Saturn V moon rocket was typical of exploration missions in that 85% of the mass lifted into orbit was propellant, as needed to leave Earth orbit, enter lunar orbit, land on the moon, return to orbit, and depart back to Earth. Reducing the cost of lofting that propellant is key to the affordability of any sustainable exploration program, at least until the infrastructure is so advanced that extraterrestrial resources can provide the needed propellant at lower effective cost. With current launch costs of ~$10,000 for every kilogram

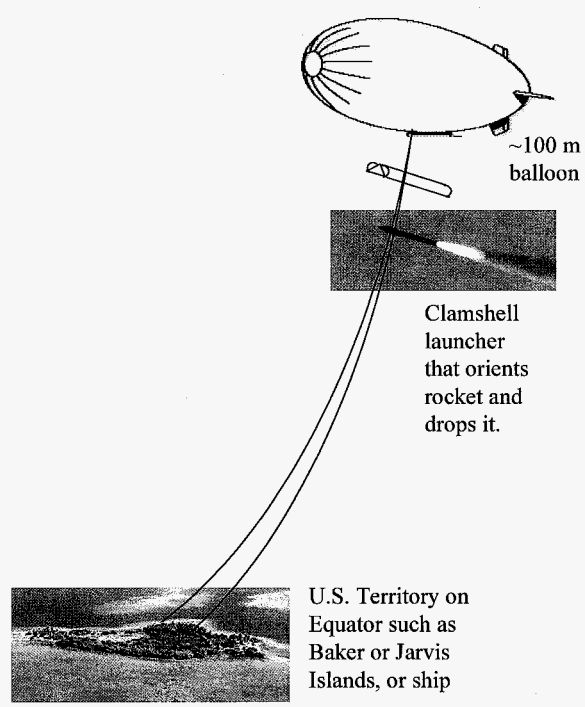

Figure 1: Concept for Low-cost Propellant Launch to LEO from a Tethered Balloon

that is delivered to LEO, and with all exploration architectures capable of captivating the public and the Congress requiring thousands of tons of propellant mass to be lofted into LEO, the cost of propellant launch can easily consume the majority of the NASA budget. The pace of exploration is almost completely limited by the rate at which propellant mass can be launched to LEO, since many exploration architectures envision reusable vehicles that can make multiple round-trips beyond LEO, so long as they have enough propellant. Thus any affordable and sustainable exploration strategy involves reduction of the cost of launching propellant by a large factor. For example, suppose that we could put dry hardware into LEO that cost (say) $100K/kg to build and launch but that could be re-used for ~25 round-trips to the moon, and we had propellant at ~$1.5K/kg available anywhere in LEO or L1 (roughly the highest cost envisioned for this concept). Each unit of hardware consumes about 4 times its own mass of propellant on each round-trip to the lunar surface and back (one to go from LEO to L1, two to go from L1 to the lunar surface and back, and one to go from L1 to LEO by propulsive injection (a low-risk but somewhat propellant-intensive approach). So a lifetime supply of propellant is 25·4·$1.5K/kg or about $150K/kg, and the total life-cycle cost of the hardware becomes $250K/kg. If NASA spent $5B/yr on this approach, with one lunar mission every month, then the sustainable amount of dry hardware that could be kept continuously moving back-and-forth through the Earth-moon system is about 40 tons. The rate of propellant use would be about 2000 tons per year. As a sanity check, one Saturn V launch every month would put about 1200 tons of propellant into LEO per year along with ~20 tons of single-use dry hardware for each mission. So this approach supports an architecture capable of a rate of exploration roughly equivalent to one "very deluxe" Apollo mission every month (e.g. having about twice the dry hardware mass) at a cost less than a third of the NASA budget (not including launch of the crew to LEO). Lofting the necessary 2000 tons of propellant by conventional means at $10K/kg would add another $20B per year, and would be completely unaffordable. Crew launch to LEO once a month need not dominate the cost of the resulting system. For example, launching the 6-ton Apollo Command Module at $10K/kg would cost only $60M per launch, assuming the module is reusable.

Reducing the cost of launch to orbit has been extensively studied.[2] Unfortunately, credible proposals to achieve significant reductions in operational costs have very large up-front investment costs. Such systems include various combinations of large reusable chemical rocket stages (with or without air-breathing 1st stages and runway launch and/or recovery), high speed "guns" that fire payloads through the atmosphere, and orbital towers or tethers that can be used as elevators to space. Advocates typically maintain that reduction of launch cost by a factor of 3 would require a capital investment at least equal to one year of the total NASA budget ($16B).[2] This is also completely unaffordable. This paper describes a means to reduce propellant launch costs (which represents ~85+% of all mass delivered to orbit in support of Exploration) by a factor of more than 15, using a capital investment that is small (~$100M), and that is demonstrable within a few years. Furthermore, in the "propellant-rich" architecture envisioned here, all other space hardware could be made significantly lower-performance (e.g. higher-mass) and hence lower-cost, but especially **more reusable** than it otherwise would, since there will be plenty of propellant available to move it beyond LEO. Thus this approach may potentially reduce the overall cost of exploration by an order of magnitude or more when extensive reuse of the dry hardware allows 95-99% of the total mass put into LEO to be low-cost propellant.

2. Small, Low-Performance, Low-Cost Rockets

The high cost of launch to orbit is not intrinsic. For example, using the rocket equation as applied to a moderate-performance three-stage solid rocket (with propellant mass fraction of 89% and specific impulse 285 s), the launch mass of a rocket will be ~33 times the mass of the payload delivered to orbital velocity. The cost of solid rocket fuel is <$13/kg[3], while moderate-performance graphite fibers for the rocket casing are available at a cost of <$11/kg[4]. If the overall rocket dry mass can be manufactured at the price of a luxury automobile ($40/kg) and the propellant can be loaded at an "installed" cost of $17/kg, then such a rocket could launch propellant to LEO with a "back of the envelope" cost estimate of ~$650/kg. This is a reduction by a factor of ~15 from current launch costs. It seems quite possible that mass production could reduce the manufacturing cost of the rockets well below this level (perhaps another factor of 2), since the relative complexity of the vehicle is much less than a luxury automobile, and the costs of solid propellant and moderate-performance graphite fibers should drop significantly when purchased in quantity.

Large rockets have been used in small quantities to reach orbit (instead of small rockets in large quantities) because of aerodynamic drag. A rocket must have a mass that is significantly larger than the mass of the atmospheric column that it displaces on its way to orbit or it will experience an unacceptable level of drag. At sea level, the atmosphere is about equivalent to a layer of 10 meters of water. So rockets that launch from sea level must be significantly more than 10 meters tall to escape the atmosphere with orbital velocity. At 26 km altitude, the atmospheric burden is equivalent to less than half a meter of water. So a small rocket can "punch through" this layer and still retain most of its velocity. Indeed, a 10-meter long rocket can launch nearly horizontally and still have acceptable drag losses.

A key problem with a small, inexpensive rocket is that each of the stages will have no thrust vector control.

Adding sophisticated guidance, and making the rocket big enough that it can launch from the surface and penetrate the atmosphere with acceptable drag losses, would make it cost as much or more than current launch vehicles. Instead, we can use a spin-stabilized rocket, and use the boiloff gasses from the cryogenic propellants that constitute the payload (LOX and either liquid hydrogen or liquid methane) to operate small "head end" vernier thrusters to counter any aerodynamic or nozzle misalignment torques on the spinning rocket. These small thrusters can also be used to intentionally precess the spinning rocket to modify the angle of attack so as to reduce drag and gravity losses, and to orient the vehicle for direct orbit injection into an appropriate circular orbit (e.g. at 280-300 km altitude). A very low cost GPS-augmented Inertial Measurement Unit[5] (IMU) and microcontroller can be used to compute the necessary thruster commands. A moderate spin rate (~70 RPM) is roughly optimal in reducing the amount of control gas that must be expended to perform accurate orbit injection. A radio receiver on the payload can be used to take commands based on precise optical tracking at the fuel depot to make the orbit corrections needed to accomplish rendezvous and to despin the payload prior to berthing. Figure 2 shows a schematic of how the rocket might be configured.

3. Launch from a Tethered Balloon

We can use these ideas to create an exceptionally simple, low-cost launch system that can be developed with very little investment. Consider a balloon that is tethered to the ground with a PBO (Zylon) high-strength rope that goes over a pulley attached to the balloon and is attached to winch drums on the ground. The balloon and pulley work like a flagpole so that a rocket and launcher can be hoisted up the tether from the ground to the base of the balloon. The balloon is stationed at an altitude above 95% or more of the atmosphere, where a relatively small rocket can be launched nearly horizontally and yet the drag losses will be low. The launcher orients the rocket using GPS and other sensing (e.g. star tracking from 26 km altitude is possible night and day) to initialize the low-cost on-board IMU and to orient the rocket to approximately the correct direction and angle, and then drops the rocket. The rocket is spun with straps around it that make it spin as it falls, like a yo-yo. Once it falls well clear of the balloon and tethers, the rocket fires. The launcher is winched back down and reloaded with another rocket. Solid-fuel or lower-pollution hybrid-fuel (solid fuel with liquid oxidizer) rockets are both attractive.

Such a system would be capable of launching rockets at frequent intervals. By launching a large number of small rockets, the same net payload can be sent to LEO as with a small number of large rockets. But a key difference is the manufacturing volume and the resulting unit cost of each rocket. A few large rockets will cost >$1000 per kilogram of dry mass, same as conventional rockets. But a very large number of small rockets can be produced with the same techniques used to mass-produce automobiles and other consumer goods. A rocket will launch from the balloon every time the fuel depot passes overhead. Launching every ~90 minutes, this is ~6000 launches per year, which is a sufficient volume to achieve most of the cost advantages of mass production. Even greater production volumes can be attained if multiple rockets are winched up to the balloon and are launched nearly simultaneously, or if multiple balloons are used. If ~250 kg of payload are launched each ~90 minutes, then a total of ~1500 metric tons of propellant would arrive at LEO each year. This is sufficient to support an ambitious and exciting space exploration program of the sort envisioned in the speech by President Bush at NASA headquarters in January 2004. This "propellant-rich" architecture allows all other hardware needed for exploration to be lower performance, somewhat heavier, more reusable, and less costly to develop than what would otherwise be practical if so much propellant were not available in LEO.

There is only one orbital plane for the fuel depot that is consistent with this volume manufacturing approach: the equator. By launching from the equator, and having the fuel depot in orbit over the equator, the fuel depot will pass over the launch site on every orbit. (Any non-equatorial launch site would pass through the orbital plane of a fuel depot only twice each day, and even then it would be rare for the fuel depot to happen to be passing

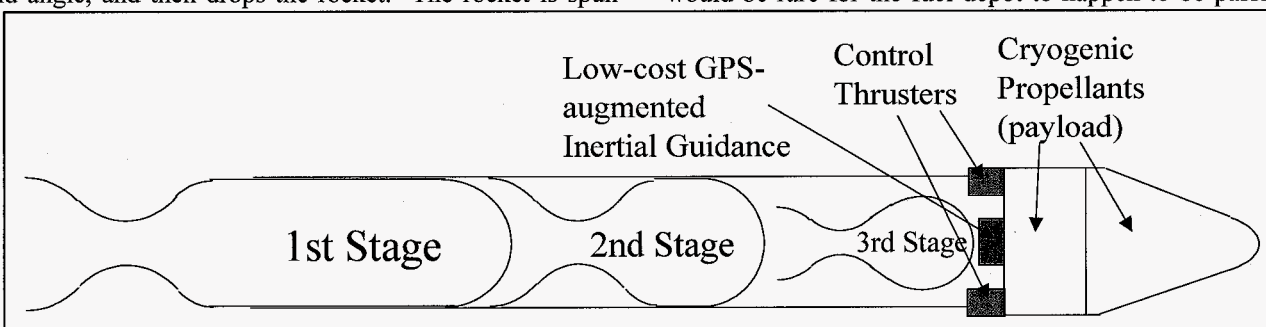

Figure 2: Rough layout of low-cost launch vehicle

over the launch site at those times as needed for direct rendezvous.) Although ship-launch from the equator is possible, it is fortunate that the US has two territories in the mid-Pacific that are within a few kilometers of the equator: Baker and Jarvis Islands. A special benefit of launching from the equator is that it is in the "doldrums," where powerful winds are rare. Indeed, because there are no Coriolis forces at the equator, there is no tendency for unstable air to organize into cyclones. Hurricanes do not occur on or cross the equator[6] and there are no jet streams[7], both of which make tethering a balloon at high altitude on the equator much easier than at other places. Another advantage of launching from extremely remote locations in the Pacific Ocean, far from normal shipping lanes, is that the most economical system will use rockets that are less than 100% reliable. The lowest cost overall system may well have a 10% launch failure rate. "3-sigma" reliability is not needed, greatly reducing the cost of every component. Commercial-grade electronics should be adequate for the control system of the propellant modules.

4. The Propellant Depot

The only place that the fuel depot can be located in order to receive frequent launches is over the equator, but in practicality this is not where the propellant is needed. Propellant is needed in other LEO inclinations (e.g. one can launch into a minimum 28.5 degree inclination from Cape Canaveral) so that vehicles in LEO can use it to boost to the moon, Mars, or beyond. Propellant is needed near the moon, where it can be loaded into lunar landers making round-trips to the lunar surface, or into Earth-moon transport vehicles that are returning to Earth. One attractive method for getting the propellant from the equatorial fuel depot to these other sites is to bundle the payload modules together, each with ~250 kg of propellant, into large "packs" that can be sent to other LEO inclinations or to the Earth-moon L1 libration point. For example, suppose that each of the payload modules has rails or clips on the side so that the robot that grapples each arriving propellant module can connect it to the previous propellant module. A long line of these payloads can be assembled. The robot would connect each payload to the previous one mechanically and also it would draw flexible cryogenic hoses out of the aft of each arriving propellant module and connect them to fittings on the previous payload. In this way, the fuel and oxidizer tanks in each propellant module can be connected in a daisy-chain to all the other fuel and oxidizer tanks (respectively) of all the other payloads in the line. The robot would probably also need to pull a multi-layer thermal blanket out of the aft of each payload and wrap it over the side and conical nose of the propellant module, reducing the amount of heat leakage into the cryogens to under 1W per square meter of exposed area[8]. Occasionally, a "special" payload module will arrive at the fuel depot. This "Hub" module will have in it a refrigerator that can re-condense boiloff gasses from the cryogenic payloads, powered by small solar arrays on the aft end of each propellant module. This module will also have six rails or clips so that six lines of payloads can be attached radially out from it. When all six lines are attached, the robot will release the assembly. The small thruster rockets on selected payloads will be commanded to spin up the whole array. This will provide enough "artificial gravity" (much less than one g) to settle the liquid to the outside of each tank, with the vapor on the inside. There will be two sets of cryogenic hoses between adjacent payload modules: one for liquid and one for vapor for both oxidizer and fuel. In this way the refrigerator can take the boiloff gasses out through the vapor lines, re-condense the vapor, and pump the liquid back into the liquid lines so that no propellant is lost.

So we imagine a "Star" configuration (Figure 3) with a "special" refrigerator and control module at the hub, and with six radial arms composed of "normal" (e.g. ultra-low cost) propellant modules. With propellant modules arriving every ~90 minutes, in about 2 weeks the robot can assemble a star having six arms each with 38 propellant modules. This "Star" would contain about 55 tons of propellant, and no propellant module would have been in orbit more than about 2 weeks. At the right time, the "special" hub module would command each of the propellant modules to fire its small thrusters to change orbit. (If one or a few of the low-cost propellant modules fail to fire correctly, that will not substantially affect the performance of the whole system, although to have reached orbit and successfully achieved rendezvous is a potent demonstration that each one has worked properly.) The Star would be spinning in the correct orientation so that the thrust from all these small rocket engines is in the direction needed either to change the LEO inclination, or to depart LEO for the Earth-moon L1 point. It takes about 3200 m/s of delta-V to leave LEO for L1, and it takes about 3800 m/s to change the inclination to 28.5 degrees. Assuming the propellant is H_2 and O_2, then leaving for L1 will consume about 62% of the propellant, while changing inclinations to match launches from Cape Canaveral will consume about 69%. This will leave about 21 tons at L1 available for refueling other vehicles, or 17 tons in LEO over Cape Canaveral, every two weeks. 17 tons of propellant is enough for a 16-ton (dry) vehicle to leave LEO for L1. If the propellant is methane (CH_4) and O_2, then leaving for L1 will consume some 73% of the propellant, and changing inclinations will consume about 80%. Despite the more adverse numbers for methane, the reduced volume of liquid methane over liquid hydrogen, and the relative ease of re-liquefying methane versus hydrogen make methane attractive. Fortunately, much more propellant is needed in L1 than is needed in LEO, and L1 is the "easier" place to get to. Remarkably, the small thrusters that are optimally-sized for controlling the spinning solid-fuel rocket as needed for precise orbit injection are also sized correctly to accomplish either the plane-change or Earth-escape in a single burn of about 15 minutes duration. The thrusters would be configured so that enough waste heat soaks

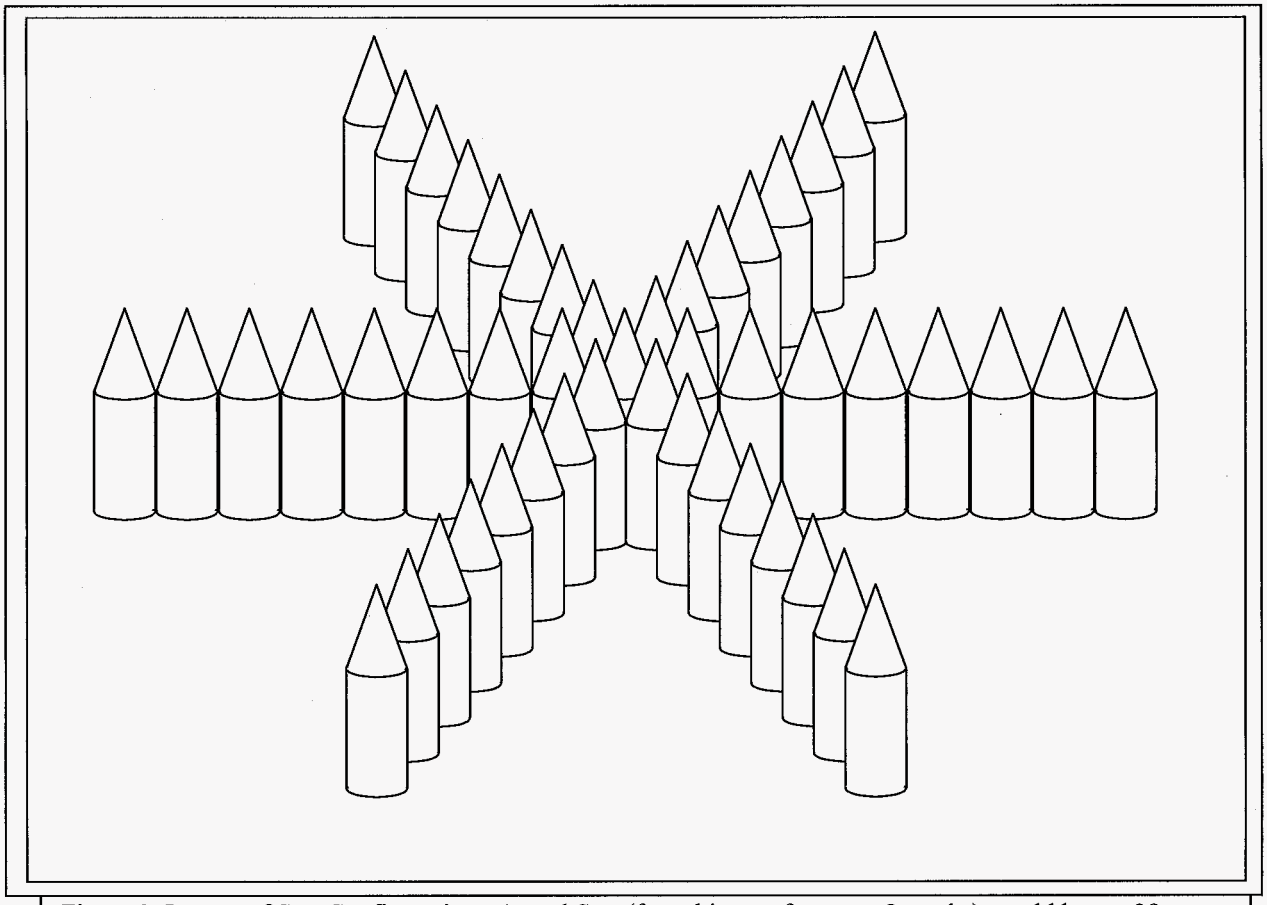

Figure 3: Layout of Star Configuration. Actual Star (for orbit transfer every 2 weeks) would have ~38 propellant modules on each arm of Star arranged around central Hub.

back to the cryogenic tanks to ensure sufficient boil-off gasses to operate the thrusters for this duration.

Once this orbit-change burn is done, the electronics, radio command receiver, and thrusters on each of the propellant modules will be "mothballed". After the spinning Star is re-oriented with it's aft toward the sun (so that the aft-end solar arrays get sustained power to run the refrigerator, and so that a reflective multilayer foil skirt at the back of each payload module protects the cryogenic tanks from direct solar impingement), a final radio command will be sent from the Hub module to fire one-shot permanent valves that prevent any leakage out of the thrusters in the event of a subsequent failure of the electronics or throttle valves. The Hub alone will be designed for a long active life, for operation through the Van Allen radiation belts and outside the Earth's protective magnetic field. The low-cost propellant module control system will not. The short-life (commercial) propellant module electronics will be powered by a non-rechargable single-use battery, with a very low-power radio receiver (like a pager) being the only element that is on continuously for the roughly two weeks between launch and the orbit transfer burn commanded by the Hub. The Hub will remain active long after the orbit transfer burn, receiving commands from Earth or other vehicles, re-condensing boil-off gasses, and cooperating with any rendezvous and re-fueling operations. It seems likely that the Star needs to keep spinning slowly throughout any fuel-transfer, since a small amount of "artificial gravity" is needed to allow all the liquid to be drained from the tanks. So the Hub will have some sort of rotating coupling on the aft central axis that allows fuel transfer while it is spinning. Once almost all of the propellant has been drained from a Star, it would be de-orbited so as to burn up over the open ocean using small control thrusters on the Hub.

This propellant-rich architecture offers many advantages. Electric propulsion is not needed (as has been proposed to lift expensive propellant out of the Earth's gravity well), nor the fuel heterogeneity or the issue of Xenon fuel scarcity that it entails. The choice of location for secondary propellant depots (e.g. Earth-Moon-L1 versus lunar orbit) is much more flexible with this architecture, since delta-V penalties are less of a concern in a propellant-rich architecture. A very small number of different rocket engine and tankage designs can be used throughout the system. But most importantly, **a total amount of propellant that supports monthly lunar missions might be launched at a cost of only about 10% of the present NASA budget, by means of a launch system whose development cost is relatively**

small. Combining this propellant-rich approach with relatively low-performance (i.e. not extensively light-weighted), rugged, and highly-reusable hardware for moving beyond LEO and the total cost of this exploration architecture becomes highly affordable. The planned Crew Exploration Vehicle, lunar landers (crew and/or cargo), and transfer stages all continue to be required in this architecture as in previous studies, although they must be configured for in-flight refueling.

5. TECHNOLOGY DEVELOPMENT

For this architecture to be effective, four technologies need to be demonstrated, which will each be discussed in turn:

- tethering of a relatively large balloon on the equator at an altitude of ~26 km, and winching rockets with mass of ~10 tons up to that balloon every ~90 minutes,
- launching a simple, small, mass-producible spin-stabilized rocket from the balloon that can deliver ~250 kg of payload to circular LEO with an injection accuracy of ~0.3%,
- rendezvous and berthing of those small payloads with a fuel depot, manipulation of those payloads (including mechanical mating of the payload modules and connection of flexible cryogenic hoses), and cryogenic fluid transfer between docked vehicles, and
- creation of the "Star Hub" cryogenic refrigerator and simple Star control system.

Tethering a balloon at high altitude: To evaluate the feasibility of launch from a tethered balloon, we have developed a simple spreadsheet model that assumes the balloon is a sphere (although it would be better if it were the shape of a blimp so that the wind drag force on the tether is reduced). Our wind model assumes a sustained wind speeds of 60 m/s (134 mph) acting on the balloon, with a drag coefficient of 0.2 (worse than a blimp, which has a drag coefficient of ~0.05), and no drag on the tether (whose total cross-section is only ~5% of the balloon). We assume the exponential scale height of the Earth's atmosphere is 7.28 km[9]. As previously mentioned, at the equator high winds are relatively rare and jet streams are non-existent, so 60 m/s winds represent a very conservative assumption. The assumed balloon fabric has a density of 100 grams per square meter, about 5 times that of typical untethered balloons that are larger than those needed for this system. This extra mass is assumed primarily to distribute the wind loads from the balloon to the tether. The assumed tether has a factor of safety of about 3 based on PBO[10]. The spreadsheet model computes the smallest system mass for each altitude, considering the competing effects that the balloon and tether have to be much larger to lift a given rocket mass to a higher altitude, but the minimum mass of rocket required to reach orbit with acceptable drag losses drops quickly as altitude is increased. The results of this spreadsheet model are shown in Figure 4. Note that the required balloon radius and mass are nearly constant at 50-60 meters and 4-5 tons, respectively, almost independent of altitude. This reflects the fact that, even though the lift of a given balloon drops rapidly with increasing altitude, the mass of the smallest rocket needed

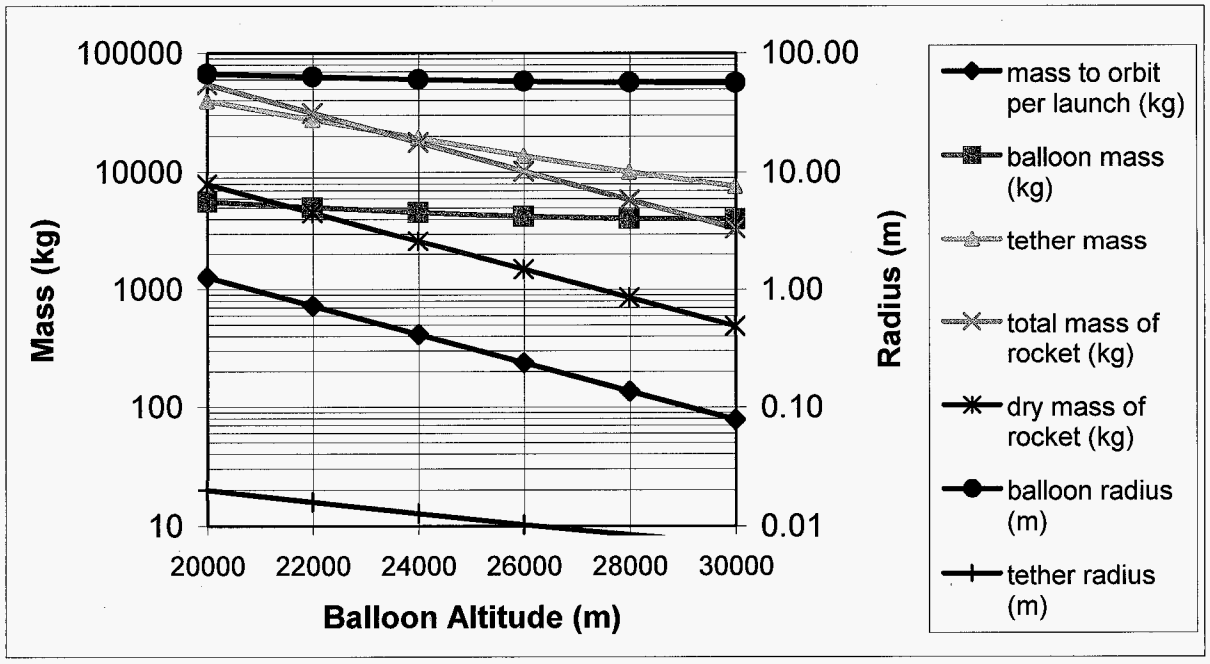

Figure 4: Results of spreadsheet model for balloon, tether, and rocket that computes a "minimum mass" system for each altitude that can launch to orbit with nominal drag losses on rocket. Note that launch of 250 kg payloads requires an altitude of 26 km and rocket mass of 10 tons.

to "punch through" the atmosphere remaining above the balloon with modest drag losses drops just as fast. Note that for payloads of 200 to 500 kg per launch, the system would operate at altitudes of 23-27 km, with a Zylon tether that is approximately one inch in diameter and whose total mass is about the same as the rocket. The total dry mass of each rocket is 1-2 tons, about the same as an automobile.

The flagpole architecture for the balloon and the pulley is not the lowest-mass approach, since the tether tension acts on both sides of the balloon pulley and so the balloon has be sized to lift twice the weight of the rocket. We have not assumed this ~2X increase in the balloon volume and tether mass in the model that generated Figure 4. However, the flagpole architecture is the simplest approach. In the flagpole configuration, there would be two winch and rocket reloading facilities - one on the north side of Baker Island, and one on the south side. The virtually-continuous prevailing surface winds at Baker Island blow from East to West at about 8 knots, in the same direction as the prevailing ocean current. Ships that re-supply Baker Island moor on the west side of the island, in the lee wind and current wake of the island where it is generally calm. The two lines that go up to the balloon in the flagpole architecture thus straddle both the island and any re-supply ships, and so there is little or no risk in the unlikely event that the tether or balloon fails and the lines fall in the water. The single tether cable is attached to one winch, goes up to the balloon, around the pulley, and down to the other winch. The total length of the tether is at least three times the distance from each winch to the balloon, so that one winch can spool the cable on one side, lowering a spent launcher from the balloon to the island, while the other launcher is simultaneously raised from the other winch as it pays out an equivalent length of line. The instant the rocket is dropped at the balloon (between the two tether lines), the loaded launcher on the other side is released and the winches start to run, so that the balloon always supports a constant payload of one fully loaded launcher and one empty launcher. (Actually each winch station would have two launcher reloading stations, so that a fully loaded launcher can be quickly attached to the tether when an empty launcher is quickly detached.) The line is winched at a speed of just over 26 km in 90 minutes, which is about 5 m/s or 12 miles per hour. Each launcher is reloaded at the surface while the other launcher orients its rocket at the balloon in preparation for launch. In a spreadsheet simulation of this flagpole system, the balloon radius would be 92 m and the tether diameter would be 4.6 cm (1.8"), using on-line performance data for commercial Zylon rope from *Yale Cordage* at a 3:1 safety factor based on the guaranteed breaking strength as needed to support the launch of 10-ton rockets from 4-ton launchers every 90 minutes at 26 km altitude in sustained winds (simulated at all altitudes) of 20 m/s (39 knots). If higher winds were experienced, no rockets would be launched, but sustained winds at all altitudes of 74 m/s (144 knots) could be tolerated at the same 3:1 safety margin. Because the two winches are widely separated, there is little danger that the two lines will become tangled (assuming the balloon is a blimp-like shape that has control surfaces to keep it under control, and the tether pulley is free to rotate about a vertical axis to keep the tethers from twisting), and each dropped rocket will fall between and well away from the two tether lines before it fires. The relatively small rockets create a relatively small shock wave at ignition whose effect is reduced by the very thin atmosphere, and they can drop many hundreds of meters or more below the balloon before ignition to mitigate any impact on the balloon or the tethers.

The first step in developing the technology for tethering a suitable balloon at the equator is to engineer, manufacture, and deploy a tethered test balloon on or near Baker or Jarvis Islands. This balloon would not have the final size and payload capacity, but it would operate at the correct altitude and have atmospheric instruments (and aircraft warnings) deployed periodically along the tether to measure the full range of atmospheric conditions over a sustained period. It would validate (or invalidate) assumptions concerning wind, temperature, and density profiles, tether and balloon fabric stress, material degradation, helium leakage, etc. (Note that in the final system there would be a small fixed tether that supports a string of aircraft warning beacons. The 200 mi exclusive economic zone around the island would be put off-limits to aircraft.) A spreadsheet simulation of such an initial testbed system has a 27-m radius balloon tethered at 26 km altitude by a 1-cm diameter Zylon rope (again using on-line data from Yale Cordage). The safety margin of this system is about 4:1 in 20 m/s (39 knot) sustained winds. The materials cost for this system based on the Yale Cordage on-line price list is about $50K. Testing of Zylon (PBO) at cold temperatures has been performed as part of other efforts, demonstrating that Zylon retains is strength and flexibility at the temperatures (-60C) associated with high altitudes.[11]

Launching a small, low-cost, mass-produced rocket to orbit from high altitude: As mentioned previously, the proposed rocket can be very simple. For example, the rocket can be spun by a cord like a yo-yo at launch, and a timer can ignite each stage at preset intervals. The first stage would ignite after a few seconds of free-fall (enough to fall clear of the launcher and tether), the second stage would light after the nominal burn time of the first stage plus a small margin, and so on. The timing is not critical for any of these events.

The nominal attitude of the rocket at ignition is about 30 degrees above the horizontal, to be determined by detailed atmospheric drag simulation. If there were no atmosphere, the rocket would be ignited horizontally, so that each stage of the rocket delivers maximum impulse in the orbital direction. However, to reduce the drag effects of whatever atmosphere remains above the balloon, the rocket is pointed somewhat upwards so that it gets out of the atmosphere more quickly than a horizontal orbit

would, at the expense of small cosine losses on each of the stages. There is a tendency for the spinning rocket to precess due to aerodynamic torques until it has reached an altitude of 60-80 km, where the air becomes so thin that the aerodynamic forces are negligible. To counter this tendency, the low-cost IMU installed in the nose of the rocket is monitored by a microcontroller to fire the "head-end" vernier control thrusters to counter any aerodynamic or nozzle misalignment torques, to manage the angle of attack to reduce drag and improve lift, and to perform an accurate direct injection into a circular orbit near the fuel depot. Final despin and rendezvous is performed using radio commands based on precise optical tracking from the fuel depot. This relatively simple and crude control system is certainly not sophisticated compared to the electronics in a luxury automobile manufactured at a comparable unit cost and production volume. In Figure 5 is shown the results of a simplified dynamic simulation of

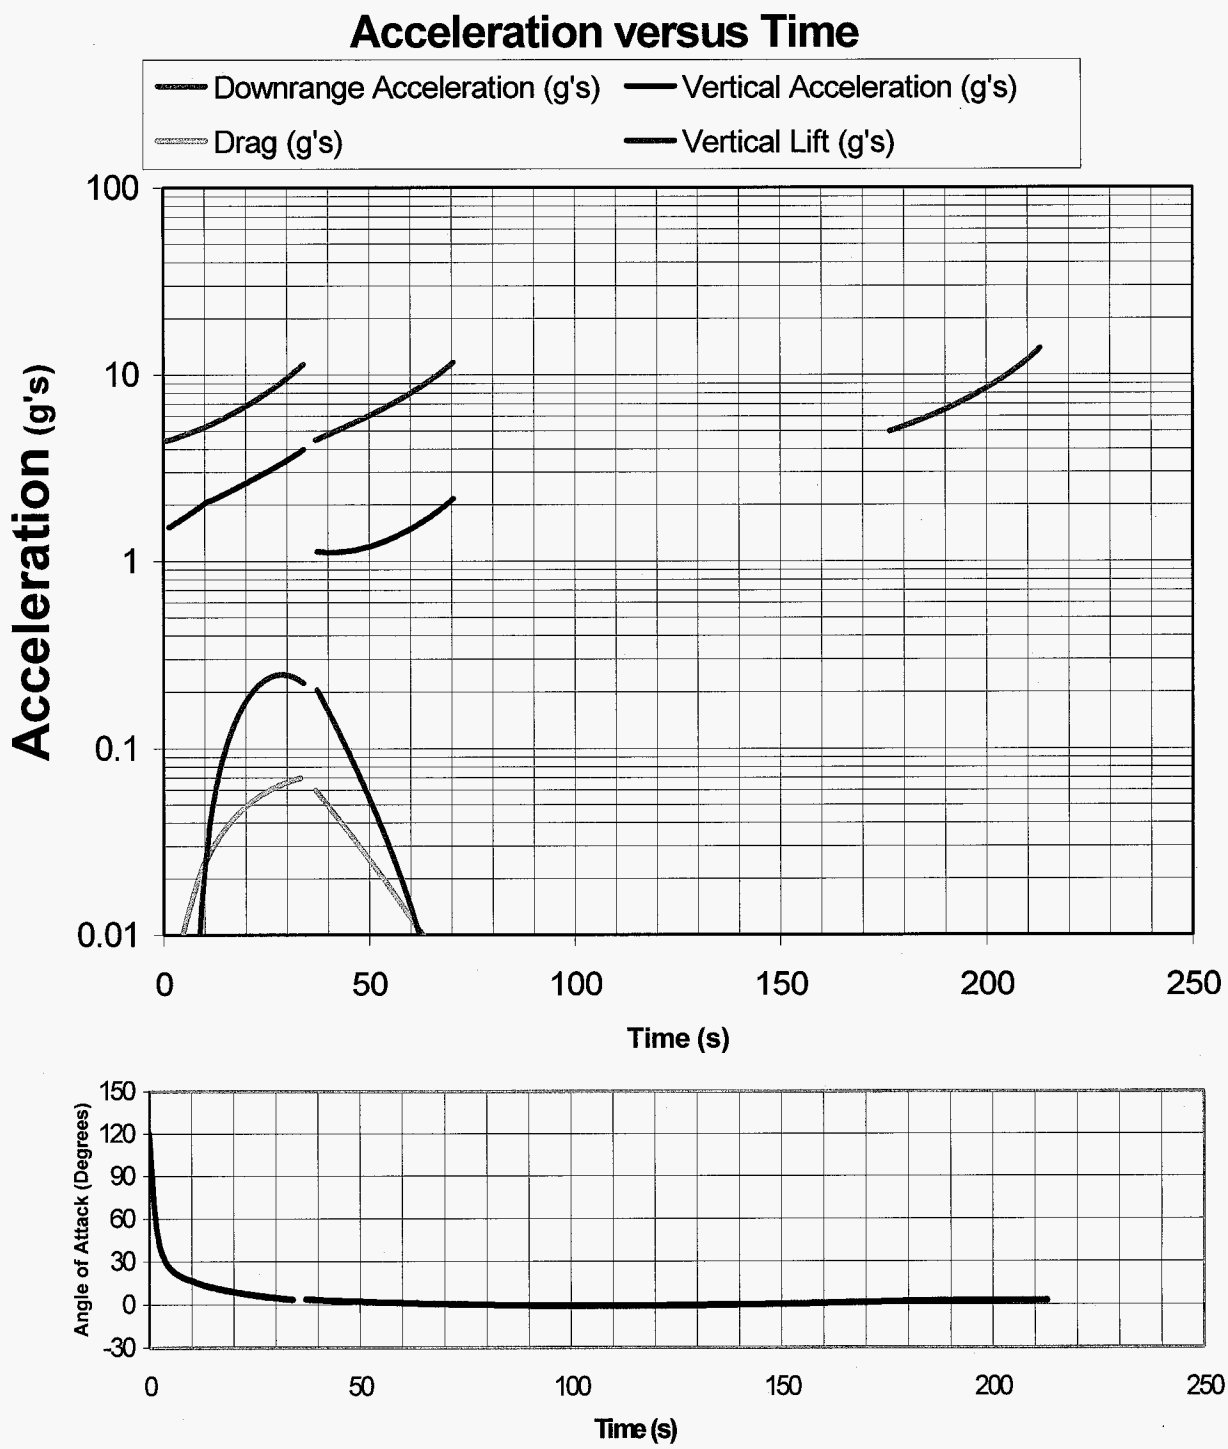

Figure 5: Simplified dynamic simulation of 10-ton, 3-stage rocket launched at 26 km, 70 RPM and 30° elevation angle

a 3-stage rocket of this type. The control system manages the aerodynamic torques and the angle of attack during the first two stage burns, and then the system coasts upward for 106 seconds to peak at an altitude of 280 km (150 nautical miles). At this altitude the control system orients the spinning assembly along the circular orbit direction, and the third stage fires, completing the injection into the fuel depot orbit. In Table 1 are the detailed parameters associated with this simulation. The optimal spin rate for this example is 69.9 RPM, balancing the increased control gas needed at high RPM to re-orient the rocket to manage the angle of attack and to align with the orbital direction against the reduced gas needed to correct for small nozzle mis-alignments (assumed ¼ degree) at high RPM.

The first step in developing the technology for the small, mass-produced rocket that can reach orbit from balloon-launch is to build a subscale version that does exactly that. Note in Figure 5 that at an altitude of 30 km or more, the mass of rocket needed to achieve orbit is only a few tons. NASA routinely launches balloons to an altitude of 32 km with a payload up to 3.6 tons[12]. A launcher and rocket, with the simple control system, would be developed for this type of untethered balloon. A test flight would be conducted. Concurrently with this demonstration program, an RFI could be sent to industry asking for expressions of interest to prototype extremely low cost rockets as needed for operation. It seems plausible that the estimates of $13/kg for solid fuel and $40/kg for casing can be further reduced by industry if they see the prospect of orders for many thousands of units per year. It would be wise to try to enlist "non-traditional" NASA suppliers to this effort, since the building of low-cost rockets is not something most NASA vendors are used to. In other industries, paying $40/kg for moderate-performance, relatively simple systems would be considered "highway robbery", while the NASA community is programmed to think of it as unrealistically cheap. It seems prudent to involve people who are part of the world of volume manufacturing outside of NASA. For example, there is a large worldwide industry making inexpensive military bombs, unguided rockets and other similar pyrotechnic devices. These industries routinely and safely handle pyrotechnic materials and produce and package them at low cost.

The Fuel Depot Robot: Rendezvous and Berthing of Payloads with the Fuel Depot. As mentioned previously, the fuel depot would track the approaching, spinning payload and issue commands via radio so that the simple control system can fire the head-end thrusters as needed to accomplish rendezvous and despin the payload. The concept is to determine the orbital parameters of the payload using observations taken from the fuel depot, and to command the firing of the control thrusters by remote control at precise times to "nudge" the payload into a rendezvous orbit. A key feature of this system is the ability to determine the precise spin phase of the payload by optical tracking, since the low-cost IMU on-board the payload will drift too far between launch and rendezvous to accurately "remember" the phase of the spinning payload in inertial space. The payloads would be painted half-white and half-black so that, when illuminated by sunlight, they appear to "blink". This blinking reveals the precise phase of rotation of the spinning payload. Knowing the orbit and spin phase, it is possible to command the firing of the small control thrusters at precise orientations in space to change the orbit as desired. Final despin would also be commanded once the payload is on a final rendezvous trajectory, leaving the payload approximately inert in space and ready for grappling and berthing.

A next step in developing and demonstrating this technology is to put a radio receiver on the payload launched from the untethered balloon in the demonstration flight described above. The orbiting payload would be optically tracked by an Earth-

Cryogenic propellant delivered to orbit	261	kg
Dry mass of payload assembly	50	kg
Dry mass of 3rd stage	88	kg
Propellant mass of 3rd stage	709	kg
Delta-V of 3rd stage	2858	m/s
Coast prior to 3rd stage ignition	106	s
Dry mass of 2nd stage	237	kg
Propellant mass of 2nd stage	1914	kg
Delta-V of 2nd stage	2474	m/s
Coast prior to 2nd stage ignition	3	s
Dry mass of 1st stage	742	kg
Propellant mass of 1st stage	6000	kg
Delta-V of 1st stage	2590	m/s
Total launch mass	10000	kg
Propellant mass fraction of solid stages	0.89	
ISP of all solid propellant	285	s
Total solid rocket delta-V	7923	m/s
Injection velocity wrt launch site	7398	m/s
Drag and gravity loss delta-V	525	m/s
Equatorial speed of launch site	463	m/s
Inertial injection velocity	7861	m/s
Injection altitude	282	km
Control gas (at optimal spin rate)	8.60	kg
Spin rate (for minimum control gas)	69.9	RPM
ISP of control thrusters	300	s
Launch elevation angle	30.0	deg
1st stage length	7.31	m
2nd stage length	2.33	m
3rd stage length	1.20	m
Diameter (payload module)	1.08	m
Drop altitude	26000	m
Drop time to 1st stage ignition	6	s

Table 1: Parameters associated with simplified dynamic simulation shown in Figure 5.

observation station and commands would be transmitted during overflights of the ground observatory to fire the control thrusters. Subsequent precise tracking of the payload over many orbits would establish that the desired trajectory modifications had been achieved and that the necessary control authority had been demonstrated. This same approach would be used to de-orbit the payload at the end of the test (although the orbit would start out intentionally short-lived in case there was a malfunction).
The cryogenic propellants in the operational system would be launched in insulated tanks. For this approach to nominally launch into the precise fuel depot orbit, the payload mass of each rocket must be tightly controlled. These tanks would need to be "topped off" in the launcher just prior to release, so that the mass is tightly controlled, independent of environmental conditions and the handling timeline. Because the payload modules are designed to be ultra-low cost, there will be no elaborate thermal control of the small cryogenic payloads. Indeed, because the control thrusters operate using the boiloff gasses from the cryogenic liquids, the payload will be designed to ensure that the boiloff rate is sufficient for adequate control during launch and rendezvous. Since about ~10 kg of total boiloff gasses are needed for precise orbit injection and rendezvous, the level of insulation of the cryogenic tanks is modest. As a result, a primary function of the fuel depot is to rapidly connect plumbing to the payload modules to connect them to a Star to prevent any further mass loss of cryogenic propellants. These functions are performed by a "Fuel Depot Robot", which is a spacecraft that has the optical tracking sensor for commanding the propellant modules to achieve rendezvous, and which also has multiple robot arms for grappling and mating the propellant modules, for attaching cryogenic hoses, and for dealing with anomalies.

The Star Hub: The Star Hub is a key component of the architecture, if it is desired not to expend the resources to develop a re-usable orbit transfer propellant tanker vehicle. It serves the function of re-liquefying any propellant that boils off, of maintaining overall control of the Star, of cooperating and assisting in the transfer of propellant to other vehicles, and in de-orbiting the Star after the propellant has been offloaded. It must survive for extended periods (perhaps 6 months or more) outside the Earth's protective magnetic field, and survive at least one transit through the Van Allen radiation belts. It has the same form and mass as a low-cost propellant module, and is launched in the same way from the tethered balloon, except that it may launch on a rocket that undergoes extra steps of inspection to provide greater reliability. Each Star Hub can cost as much as a few million dollars without adversely affecting the overall economics of the system, since there is only one Star Hub for every 228 propellant modules (which each cost ~$160K including launch, so the whole Star has a value of over $36M). Since the Star Hub has a mass of ~300 kg, a cost per kilogram of up to $10K/kg is acceptable, and is entirely consistent with other sophisticated aerospace hardware that is manufactured in quantities of ~100, such as the Iridium satellite constellation. The likely development cost of the Star Hub is a few hundred million dollars. One issue is that it may be required to attach small active thruster pods on the ends of each of the arms of the Star to give sufficient control authority to precess the spinning star into new orientations. This is because torques generated by thrusters on the Hub may not have enough moment arm to adequately torque the Star. If more detailed analysis indicates that this is the case, then presumably long-life, radio controlled, solar powered thruster pods would be developed that can be attached by the Fuel Depot Robot at the ends of each arm of the Star. These thruster pods would attach to the last propellant module in the daisy chain, and thus have access to the full propellant supply. The thruster pods would communicate using the same wireless protocol that is used to communicate between the Star Hub (and the Fuel Depot Robot) and the low-cost propellant modules prior to their being mothballed.

The sequence of operations envisioned is that every two weeks a new Star Hub would be launched into orbit. It would rendezvous with the Fuel Depot Robot in the same fashion as a propellant module (using a small amount of propellant that "primes" the refrigerator). The Robot would take the small thruster pods off the Hub and store them in receptacles on the Robot. Ninety minutes later, the Fuel depot robot will hold onto the Hub while it makes the rendezvous with the first propellant module that will be part of the new Star. Once the first propellant module is mated to the Hub, with the robot hanging on to both of them, the thrusters on the propellant module will be commanded to fire to re-boost the whole assembly to a slightly higher orbit, to make up for whatever orbit decay has occurred during the fabrication of the previous Star. Every 90 minutes additional propellant modules would be mated to the previous one, until one "arm" of the Star is complete (38 propellant modules). One of the thruster pods would be affixed by the Robot to the end of the arm, and the cryogenic hoses attached to the last propellant module on the end of the arm. Then the Robot would climb down the arm until it had reached the Hub. There it would start building a new arm with the next propellant module that is launched. This process would repeat until all six arms of the Star are complete. The Robot would push off from the Star. Once the Star had drifted a safe distance away, the Hub would command the thrusters to orient and spin the Star into the correct attitude for the orbit transfer burn. At the proper time, this burn would commence, and the Star would leave the fuel depot orbit.

Robotics is a key element of this concept. Performing routine tasks in space such as berthing the small payloads every 90 minutes, cryogenic refueling of docked spacecraft, or assembling the Star from the Hub and multiple propellant modules is made affordable by using robots supervised by humans at a ground control station. It is unlikely that humans will ever go to the fuel depot in equatorial LEO since there is no human-rated launch site

on the equator. Because the cost of a human in LEO is ~$100K/hr, while the cost of a human on the ground is ~$100/hr, this otherwise large cost is reduced by a factor of 1000 by using a ground-supervised telerobot. The on-board capability required of the robot is "just" the act of being able to apply the proper forces within safe bounds as needed to accomplish the task, while "autonomy" and "intelligence" on the robot (as opposed to residing in the human on the ground) is enhancing but not enabling.

Of course being able to apply the proper forces is not trivial. The simplest robot suitable for this mission needs to have sufficient mechanical degrees of freedom to physically accomplish all of the tasks, and needs to possess enough force capacity. It needs accurate position and force sensors. It needs an appropriate set of position and force servo loops. And it needs to have sufficient sensors (cameras, etc.) for the human (at least) to get adequate situation awareness. But that is about all it needs.

Consider a system where the style of interaction by the human operator is for the human operator to say "put that there", "insert that screw into that hole", "pick up that cryogenic hose connector and screw it onto that mating connector", and the like, using appropriate cursor designations on stereo imagery to bind the pronouns with things the robot can understand. Simple machine vision algorithms can be used to determine the orientation of the objects once they are designated, and the machine can query the human for additional guidance if it gets confused. An effective style of interaction between the human and robot is analogous to an engineer directing a competent 12-year-old to build a carburetor: everything needs to be directed in considerable detail but the job can get done fairly quickly. If the child knows some primitive skills like how to insert a screw without getting it cross-threaded, then this sequence does not need to be described in excruciating detail every time. Other useful skills might include how to affix velcro fasteners on a thermal blanket or how to connect two payload modules to each other using pre-installed "male" and "female" rails, etc.

This author published an IEEE paper[13] "An Evolutionary Strategy for Telerobotic Manipulation via Geosychronous Relay" in 1989, which introduced the concept of "mechanical primitives" for time-delayed manipulation. All tasks can be performed using a mechanical primitive of the form denoted by this PI as "slide until." The end-effector or workpiece always has some number N degrees-of-freedom constrained by sliding task contact, with the other "6 minus N" DOFs being unconstrained. In limiting cases there could be no contact and so N=0, or the object is frozen in place and N=6. With each primitive motion, the robot will move the component part through some trajectory in the (6-N)-dimensional free space, while maintaining sliding motion in the N-dimensional "constrained space". This sliding motion continues "until" some termination condition is met, such as a sudden increase (or reduction) in force or torque.

An example will make this more clear. A sequence for inserting a machine screw into a threaded hole in a plate might be as follows:

1) Move the tip of the screw until it contacts the surface of the plate. (A free-body move is a degenerate "slide" where N=0, that is, no degrees of freedom are constrained by contact.)
2) Slide the tip of the screw along the plate until it drops into the threaded hole. This slide is terminated when the momentary loss of surface contact is detected as the tip falls into the hole.
3) Rotate (a rotational form of "slide") the screw about the tip in two dimensions until the screw axis is nominally aligned with the expected axis of the threaded hole.
4) Rotate the screw backwards (counterclockwise for right-hand threads) around the nominal screw axis until the threads "snap" over each other, indicating that the threads are ready to engage.
5) Rotate the screw forwards (clockwise) until the desired torque is reached after a threshold number of turns. If the running torque rises above a low threshold before the anticipated number of turns, then cross-threading is suspected; alert operator or return to step 4.

In each case, the action is to "slide" the object nominally in a free direction, maintaining modest contact forces in each of the N orthogonal constrained directions, "until" some termination condition is met (e.g. an increase in force or torque, a reduction in force or torque, a "clicking" is felt, etc.) All tasks can be decomposed into a sequence of such "slide until" primitives. The common sequences, such as inserting a threaded fastener, opening a drawer or door, etc. can be stored and used as "macros."

Simple machine vision algorithms that can recognize and track pre-defined and designated handles, tools, fasteners, hinges etc. are straightforward with existing technology and will be very helpful, although an alternative and fallback is for the human operator to call up wire-frame overlays so the human can designate them on the stereo images taken at the worksite to tell the robot explicitly the position and orientation of each object to be contacted, manipulated or avoided.

Such a system that combines the remote-site robot and human control station to perform all the needed tasks can be demonstrated on the ground for a modest cost. We propose to perform such a demonstration as an initial part of this effort. It will "prove" that this simple style of control and manipulation is effective despite time delays of several seconds (e.g. anywhere in Earth orbit). A skillful human operator should be able to keep the command buffer filled with commands of the style "move that there, move this here..." so that the motion of the robot remains continuous at the limiting speed set by task safety protocols. Whole-arm collision avoidance and end-effector force thresholds can be set conservatively so that

any anomalies cause the robot to stop and alert the human without damage to any part of the task. The human will watch the task execution via high-resolution stereo imagery and will have sufficient situational awareness to replan and recover from anomalies.

Perhaps the most important part of this supervisory control concept is that it is achievable immediately, and does not require any of the elusive and ill-defined "autonomy" or "machine intelligence" functions that have been oversold and underachieved so much in the past that it has to a large degree discredited the whole field of robotics. Whatever such techniques become available in the future will of course be enhancing, but the huge cost savings associated with getting the human out of the space suit and putting them on the ground will already be achieved.

Such a supervised robot will be capable of assembling each Star from a Hub module and many propellant modules, and will be the only element that permanently resides in the fuel depot equatorial orbit.

5. Conclusions

Reduction of launch cost for propellant by a factor of 15 to 30 below current costs of ~$10,000 per kilogram appears to be possible by using small, simple, mass-produced rockets launched from a high-altitude tethered balloon. It also appears that such a system can be developed with a relatively low capital investment, compared to alternative means that have been proposed to reduce launch costs. This would allow NASA to adopt a "propellant-rich" architecture that would eliminate the need for electric propulsion tugs to move propellant or cargo from LEO to high orbits, would allow all hardware to be launch dry, reducing or eliminating the need for a new heavy-lift launch vehicle, would allow all the dry hardware to be more massive, lower-performance, lower-cost, safer and more reusable than would otherwise be possible, and would reduce any near-term need to produce propellant from lunar, Mars, or asteroidal resources. The overall effect could be to reduce the cost of implementing the Exploration Vision by an order of magnitude.

Aknowledgement

The research described in this publication was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

Appendix A: Notes on the "Simplified" Launch Simulation of Figure 5

• The vertical acceleration is plotted after subtracting gravity (e.g. a vertical thrust equal to the weight produces zero acceleration). It is "simplified" in that nutation effects are not modeled.

• Small tabs are presumed to be attached to the stages to trim the Center of Pressure so that the CP has approximately the correct relationship to the Center of Mass during the burn of each stage. The CM starts behind the CP, but moves forward as each stage burns, ending ahead of the CP. This reduces the amount of control gas required to maintain attitude control.

• The lift and drag coefficients as a function of angle of attack are cubic and quadratic fits (respectively) to the wind tunnel measurements for a cylinder with a conical nose at Mach 3.2 as published in http://naca.larc.nasa.gov/digidoc/report/rm/11/NACA-RM-E53L11B.PDF.

• Maximum aerodynamic flow power density is 2.73 W/cm^2 at 34.8 km elevation and 2227 m/s velocity. Heating is > $1 W/cm^2$ for 30.9 s. Total aerodynamic flow energy density 74 J/cm^2, which would boil off a total of ~3 kg of cyrogenic propellants (if the air delivered 100% of kinetic energy over 1.1 m^2 frontal area).

Appendix B: Frequently Asked Questions

Q: Why is propellant 85-95+% of all the mass needed in LEO to support the Exploration Vision? A: The Saturn V moon rocket is a good example of an exploration system, with about 85% of the mass lofted to LEO by that system being propellant. Even with the highest-performance cryogenic propellants, over half of the mass of any spacecraft leaving LEO for the moon or beyond has to be propellant just to escape Earth's gravity. If that spacecraft is also going to enter orbit or land on another body, more propellant is needed. If, unlike Apollo, much or most of the hardware can be reused, then the 15% of the Apollo mass lofted to LEO that was dry hardware will drop to 5% or less with as few as 3 re-uses of each piece of hardware. If the hardware can be used 10 or more times, the percentage of the total mass that is propellant needed in LEO to support exploration could approach or exceed 99%.

Q: Balloon-launch isn't new. What is new here? A: This concept combines two ideas from the 1950s: balloon-launch[14], and spin-stabilized "unguided" launch to orbit at low inclination from high altitude[15-20], to potentially reduce the cost of the NASA exploration program by an order of magnitude. An enabling breakthrough is the development of PBO (Zylon) rope, which can support its own suspended weight from a height of over 100 km[10]. There are concerns that Zylon may degrade over a few years, but because the cost of the tether is not very much greater than the cost of a single launch, the tether (and the balloon) can be replaced every year or so without affecting the overall economics of the approach. This enables the key idea: to reduce the cost of launchers from >$1000/kg (dry) to ~$40/kg by mass-producing very simple, low-tech rockets and launching them above most of the atmosphere. These rockets are simple and low-tech because they are basically just shells

filled with pyrotechnic materials, with only the simplest of guidance, sensors, and actuators. The off-gassing of the cryogenic payload provides all the control authority needed to maintain the proper angle of attack during orbit injection, and to perform rendezvous, while the remaining liquid represents 85-95% of all the mass that NASA will need to launch into orbit. So we can reduce the size of the rockets being developed, and increase their production volume, to configure the most economical system. Practically, we can winch one or a few rockets up the tether every time a fuel depot in equatorial orbit passes over the launch site, and so can launch any multiple of ~6000 rockets per year. This production volume is comparable to luxury automobiles that have many more manufacturing steps and quality requirements but cost ~$40/kg.

Q: Why only ~250 kg of propellant per launch? A: What is important is to get close to the lowest possible launch cost in $/kg, so the rockets need to be manufactured in the largest possible volume. Dividing a minimum desired propellant launch rate of 1500 tons per year by the number of launches (5844 at once-per-orbit) gives 273 kg of propellant per launch. If two rockets are launched on each pass of the fuel depot, then ~3000 tons per year could be put into LEO. The size of rocket needed for this payload implies a balloon and tether size that is comfortably within the range of operational experience, and will have low drag losses when launched from ~26 km altitude.

Q: Why launch from ~26 km? A: This has not yet been fully optimized. However, the spreadsheet model of Figure 4 includes the following effects: the increased balloon and tether size needed at higher altitudes, the declining atmospheric drag losses at higher altitudes, the time to hoist the launcher up the tether between launches, and winds. It appears that subsequent detailed atmospheric simulation will produce a launch cost minimum at about 26 km (85,000 feet) for payloads of ~250 kg each.

Q: Will the balloon survive for long? A: Yes. The spreadsheet model referred-to previously assumes a bag fabric with a mass of $0.1 kg/m^2$, which is about 5 times that used in most large high-altitude balloons. The tether is assumed to be Zylon (PBO) with a 3:1 safety margin[10]. The rate of helium leakage out of the bag has not been calculated, but the cost of helium replenishment is sure to be negligible compared to the ~$160K cost of the rocket that is launched every 90 minutes. The shock of rocket ignition on the bag has not been calculated, but it would have almost no impact on the overall system cost if the rocket needed to be dropped ~1km below the balloon prior to ignition. Note that the rocket is relatively small (~10 tons) compared to normal Expendable Launch Vehicles (~800 tons), and so produces a relatively small shock. Also note that shock coupling at high altitude is greatly reduced compared to sea level.

Q: Do you really use the balloon just like a flagpole? A: Yes. There would be two rocket reloading stations: one at the northern-most point of Baker or Jarvis Islands and one at the southern-most point. The tether would circulate between these two stations, cycling over a pulley at the base of the balloon, just like a flagpole. Winches at the two stations would circulate the tether. There would be two rocket launchers, one of which would be loading at one station while the other is launching at the balloon. The required winch speed is ~ 5 meters/second (12 mph). A small, stationary Zylon tether would hold a string of warning beacons and radar reflectors to warn the occasional aircraft that strays into the area. Because the prevailing winds blow almost constantly (at about 8 knots) from East to West, the balloon will not actually sit over the island, and if there is any failure of the balloon, tether, or launcher all parts will fall into the ocean. The two tether lines will form a "V" at the balloon, so the rocket will fall between the two tether lines and not put either one at risk during ignition. During the rare periods of high winds no launches would occur, and the tether will have substantial margin even in Category-5 hurricane-force winds.

Q: How do you point the rocket as it is dropped? A: The launcher assembly will torque the rocket against the tether line, reaction wheels, and using aerodynamic control surfaces to point the rocket in approximately the desired orientation, and to initialize the low-cost IMU, as measured by GPS and other sensing (e.g. star tracking from 26 km altitude is possible night and day). The selected attitude of the rocket will be some 30 degrees above horizontal, to be determined by detailed atmospheric drag simulation. If there were no atmosphere, the rocket would be ignited horizontally so that each stage of the rocket has maximum impulse in the orbital direction. However, to reduce the drag effects of the atmosphere, the rocket is pointed slightly upwards so that it gets out of the atmosphere more quickly than a horizontal orbit would, at the expense of small cosine losses on each of the stages. The acceleration of the rocket needs to be high enough that the vertical component of the acceleration is greater than gravity, and the lift and "curve ball" forces on the spinning rocket need to be fully accounted-for in the detailed aerodynamic simulation.

Q: Can you launch larger payloads of other types? A: Yes, in principle, but to a large degree that misses the point of this concept. The purpose of this system is to reduce the cost of propellant launch, which represents 85-95+% of all the mass needed in LEO to support the full vision of exploration. Most other mass sent to LEO will be expensive, reusable infrastructure that is worth launching by the most reliable conventional techniques. Most piloted and high-value cargo launches from the US will presumably occur from Kennedy Space Center, not once-per-orbit from the equator as envisioned here. If there were some other architectural element that

represented a huge total mass, and could be launched frequently from the equator in relatively small chunks, then it would be suitable for this approach. Developing a very large balloon, a very thick tether, a very large launcher, and a very large rocket so that larger payloads of other sorts could occasionally be launched would not be economical compared to existing launch vehicles.

Q: What about pollution? A: This could become a concern over time. Eventually the solid rockets can be replaced with cleaner hybrid rockets (LOX or N_2O liquid oxidizer and rubber or PMMA (Lucite/Plexiglas) fuel) at a modest economic penalty. Replacing the 1st stage alone would account for over 60% of the total propellant.

Q: What are the safety implications? A: The proposed launch sites are far from the normal shipping and commercial airline corridors, and are U.S. territories that have many hundreds of miles of open-ocean Exclusive Economic Zone all around that can be put "off limits" to ships and aircraft. The first two stages would fall into the sea within this restricted range. Warning lights for aircraft would be put on the balloon and hung on a stationary tether. The lights in the tether would be fiber-optic, so as not to require electrical conductors that would become giant lightning rods. It can be arranged so that the 3rd stage burns up over the ocean after a partial orbit. This is accomplished by breaking the 3rd stage of Figure 5/Table 1 into two stages that fire in rapid succession - a 3rd stage with approximately 2750 m/s of delta-V, and a 4th stage with the last ~100 m/s of delta-V needed for final orbit injection. Because the spent 3rd stage won't achieve orbital velocity, it will burn up in the atmosphere after less than half an orbit - in the central Atlantic between South America and Africa. The very small 4th stage will remain attached to the payload and be disposed of in the same fashion as the rest of the empty payload shell (see below for details). The use of additional propellant boil-off gasses can substitute for the use of a 4th stage.

Q: Tethering a high altitude balloon is difficult. Aren't you glossing over the problems? A: There certainly are problems that need to be addressed, but the equator is probably the easiest place on Earth to tether a balloon at high altitude. The weather and winds on the equator are much more benign than anywhere else. These are the "doldrums" that bedeviled ancient sailors. The overall circulation pattern of the atmosphere is composed of three circulation cells in each hemisphere - tropical, temperate, and arctic. Two jet streams circle the planet in each hemisphere at the boundaries of the temperate cells. Generally at the equator the atmosphere is in a gentle updraft before it enters one of the adjacent circulation cells, there is no jet stream, and there is no Coriolis force to organize unstable air into cyclones. Large storms do not cross the equator because of the lack of Coriolis force. Organizations such as TCOM are working to show that high-altitude tethered balloons (called aerostats) are useful for telecommunications relay, homeland defense, and other applications. They have published extensively[21]. Deploying aerostats through a jet stream is a major problem, but is not a concern here. Lightning is an issue, although less so on the equator due to reduced storm activity. Degradation of the tether or balloon material due to UV and exposure is an issue that applies to equatorial tethering, but as mentioned previously these items can be replaced regularly without seriously affecting the overall economics of the concept.

Q: Isn't the Defense Department working on this? A: Balloon launch and high-altitude launch of rockets dates to the 1950's for the fundamental reason given previously: it reduces the mass of the atmospheric column displaced by the rocket on its way to space. Thus this same approach holds promise using conventional precision-guided rockets to "flatten" the launch cost curve, where today the launch cost per kilogram explodes in the case of very small payloads. The rumored DOD program presumably addresses this, using precision-guided vehicles, for launch from the continental US where the jet stream and frequent storms pose severe development challenges. There is no present DOD program to launch propellant into orbit for operations beyond LEO.

Q: How do the "head-end, vernier" thrusters work? A: These thrusters burn the boil-off gasses from the cryogenic propellants that make up the payload, supplied under the natural pressure maintained by the boil-off process. Small valves are actuated by the microcontroller (either by monitoring the IMU or by radio command from the fuel depot) to admit the propellant gasses to each combustion chamber, along with a source of ignition (e.g. a spark plug). The small rockets have a combustion rate of a few tens of grams of propellant per second, adequate to each burn ~5 kilograms of propellant during the ~250 second interval from launch to orbit insertion (see Figure 5). These rockets do not need to have a long service life (just over a thousand seconds), and can use an expendable nozzle throat (e.g. graphite), much like a solid rocket. However, it is probably best if they are gimbaled so that as few as two thrusters suffice for all necessary functions. With two gimbaled thrusters at the outside rear of the payload (see Figure 2), they can thrust mostly outward to provide "head-end" vernier control of the larger vehicle during launch and orbit injection or to re-orient the payload during rendezvous, or can thrust through the center of mass of the payload to accelerate the payload either down-track or cross-track (as determined by the payload orientation) as needed to accomplish rendezvous, or can thrust nearly tangentially to de-spin the payload prior to capture by robotic arms for berthing at the fuel depot. The two gimbaled thrusters, each with two gas valves and two gimbal actuators, probably form the simplest control actuation scheme that can accomplish all the needed functions. It may be desirable to arrange for some or most of the heat from each thruster to leak back into the cryogenic tanks to sustain the flow of boil-off

gasses during long burns. This will also cool the thrusters to prolong their life and perhaps reduce their cost. In the simulation of Figure 5 and Table 1 the ISP of the control thrusters is assumed to be only 300 s, but in fact it should be possible to achieve 350 to 450 s using cryogenic boil-off gasses. The impulse uncertainty of a solid rocket motor is typically 0.3% (see the ATK-Thiokol "Space Rocket Motor" catalog), so one would expect that the total amount of control propellant needed to rendezvous (even if the attitude were always perfectly controlled) would be 0.3% of the total solid propellant load (about 8600 kg) or 25 kg. In fact, the control system could modify the angle of attack during the second-stage burn to account for the impulse uncertainty of the first stage (which represents 70% of the total solid propellant). If the first stage delivered excessive impulse (beyond the minimum required) then the control system could intentionally let the vehicle precess around the local aerodynamic velocity vector to "burn up" the excess speed by allowing the drag to be larger than it could be, and to allow the lift direction to cycle between upwards and downwards. As soon as the excess velocity is lost, the control system would stop the precession and reduce the angle of attack to normal. This will reduce the amount of control gas needed to achieve rendezvous to about 10 kg (in addition to the 8 kg needed during orbit insertion).

Q: Is this concept related to orbital towers or tethers that can be used as an "elevator to space"? A: No. Those concepts require incredibly large up-front investments, using super-strong materials that do not yet exist that can support their own weight when suspended thousands of kilometers in Earth's gravity. While it is true that our concept is enabled by the commercial availability of PBO (Zylon) rope that can support its own weight when suspended from altitudes of over 100 km, our approach winches the rocket a negligible fraction of the way out of the Earth's gravity well (only 0.4%) before the rocket is fired. The purpose of our tether is to winch the small rockets above enough of the atmosphere just to the point where the drag on a small rocket is acceptable, so we can launch frequently and enjoy the benefits of mass production. We are not attempting to winch payloads all the way out of the atmosphere into space; instead we are merely taking advantage of the fact that the rocket equation does not contain any factors that relate to the absolute scale of the system, so many small rockets can reach orbit as well as a few large ones if atmospheric drag can be neglected. Those many small rockets can be mass-produced, and don't require a large up-front investment.

Q: Wouldn't it be better to launch the cryogenic payloads directly to an escape orbit toward L1, instead of just going to LEO? A: No. It is true that the temperature of a payload in radiative equilibrium (behind a sun-shield) drops as one departs from the Earth, to the point where it is below the boiling point of LOX at a distance of about 4 Earth radii above the surface. So the boiloff of the cryogenic payload could be reduced or eliminated by promptly getting away from the Earth. However, the cost implications on the propellant modules are too great to justify this approach. Making a propellant module that can operate for up to two weeks in equatorial LEO is very different from making a payload module that can operate for months, operating through the Van Allen radiation belts and outside the protective cocoon of the Earth's magnetic field and thermal radiation. The moon's orbit is significantly inclined with respect to the equator, so one can really only launch to L1 from the equator twice a month with minimum propellant use (timed so that L1 crosses the plane of the equator at the time of arrival). In principle, any terrestrial launch site can launch at any time to any specified point far from the Earth, since the launch site, the center of the Earth, and any distant point define a unique plane. The vehicle would launch into this plane, and then make a partial orbit in LEO until it was at the right place for another burn to start a parabolic trajectory to the distant point. However, the necessary midcourse corrections and final rendezvous at this remote fuel depot would require that the electronics, sensors, valves, actuators and thrusters on the propellant module remain functional throughout this long journey outside the Earth's protective magnetic field. To prevent boil-off of the precious cargo, the entire assembly would need to drop to the cryogenic temperatures of deep space. As a result, the only launch and fuel depot sites that allow frequent launch and arrival (e.g. thousands per year) without requiring extended-life, radiation-hard, and extreme-temperature electronics and actuators on the payload are on the equator and in purely-equatorial LEO, respectively. A side benefit is that the rockets always launch due East, taking maximum advantage of the Earth's rotation, getting a solid fuel savings of ~18% over the worst-case (polar) launch from an arbitrary site into LEO as sometimes needed for direct launch to L1.

Q: How do you arrange the plumbing so that you can get all the liquid out of the tanks? A: There are many possible configurations for the cryogenic tanks to be arranged in the "cylinder with cone" payload volume. During launch, it is spinning around the symmetry axis, and it is accelerating toward the tip of the cone, so it is natural to put the vapor intake near the central axis and up toward the tip of the cone. It is also natural to put the liquid intake near the base of the tank, along the outside. When the propellant module is spinning as part of the Star, the centrifugal force will cause the liquid to sit on one side of the tank, with the vapor on the other. So if the liquid intake is arranged at the base and outside of the tank at the point closest to the attachment fittings on the outside of the Star, then it will serve to collect liquid in either the launch or the Star configuration. It may be that two vapor inlets are needed - one as close to the apex of the cone as possible, and one as close to the inside of the Star as possible. It may be that a simple float valve could close off either vapor inlet if it is immersed in liquid. When the propellant module is neither spinning nor thrusting, there is no way to separate the liquid from the

vapor. The two sets of cyrogenic hoses will be arranged so that all the liquid inlets are "T-ed" together, and all the vapor inlets are also "T-ed" together, so that the Star Hub can extract vapor from the vapor hose line (under the small pressure generated by the boil-off), re-condense it, and inject it into the liquid hose line. When the Star is spinning, all the liquid will drain to the outermost modules on each arm of the Star, while the innermost modules may contain nothing but vapor. The vapor pressure will be maintained at a level no greater than that needed to push the liquid "uphill" to the hub against the centrifugal force of the spinning Star. Since the Star will spin very slowly, this centrifugal force will be much, much less than the equivalent of one Earth "g".

Q: Will it be possible to build an adequate cryogenic refrigerator into the volume and mass of a single payload module? A: In reference [8] it shows that multilayer insulation blankets are capable of reducing the heat flux between room temperature and liquid nitrogen temperature to under 1 W/m^2. The spinning Star will be normally oriented with the aft pointing toward the sun, where a reflective shield will protect against solar impingement. Thus the front and sides of each payload module will "see" the heat radiated by the Earth about half the time, and will "see" deep space the other half of the time. So it is prudent to assume that the heat leak will be about 0.5 W/m^2 over the exposed surface of the Star. Each payload will be about 1 m in diameter and 2 m tall, so the exposed area of each propellant module in the Star configuration will be about 4 m^2 (a little under 2 on each side, because of the nose cone, plus the back). With $38 \cdot 6 = 228$ total propellant modules, the combined heat leak will be about 500 W. The efficiency of a Carnot cylce (e.g. Sterling refrigerator based on helium gas) operating between (say) 40K and 400K will be only about 10%, so about 5KW of total power will be required. If there were a solar panel on the back of each propellant module that generated 20W, then the total amount of power generated by all the panels would be adequate. Presumably, all panels on any one arm of the Star would be in parallel, to be "fail-safe" against any single failure (with diode isolation), while the six arms can be connected in series to get higher voltages for more efficient operation of the large refrigerator. This would require that the robot make an electrical as well as the mechanical and plumbing connections between the propellant modules. A low-cost silicon solar array (15% efficient) would only need a cross-section of about 10% of the total cross-section of each payload (1.1 m^2) to generate the needed power. Presumably the robot would attach a 5KW solar array to each hub when it starts to build a star, and disconnect it when the star is complete. The robot would retain the 5KW solar array for use in building the next star. There are many possible configurations of the tanks in the low-cost payloads, presumably all based on concentric cylindrical walls inside the "cylinder with cone" form-factor of the payload. If the fuel is H_2 (instead of methane) then presumably the inner cylinder will hold the hydrogen, and there will be multilayer insulation between the walls of the hydrogen and oxygen tanks. This is because liquid hydrogen is so much colder than liquid oxygen, and the difficulty of re-refrigerating liquid hydrogen is so great, that the heat leak into the hydrogen tank must be carefully controlled. If it turns out that the mass and size of the Hub module needed to support 228 low-cost propellant modules is too great, then instead the Hub will presumably be designed to support half that many (e.g. one week's worth of launches), or some integer divisor. The importance of making the Star in modular sub-units of two-weeks production is that a minimum-energy launch from equatorial LEO to L1 is possible every two weeks (e.g. when the moon passes through the plane of the equator), and most of the propellant must go to L1. At L1 the "radiative equilibrium temperature" behind a sun-shield is below 40K, so very little refrigeration is needed.

Q: How do you do a piloted Mars mission with this architecture? A: The safest way to do a piloted Mars mission is to have a single stack of components to perform the Mars mission that allows some basic return-to-Earth capability at any time in the event of a major malfunction. This stack would have, at a minimum, sufficient propellant to leave Earth orbit on a Trans-Mars trajectory, to enter Mars orbit, to leave Mars orbit, and to enter a Trans-Earth trajectory from Mars. Since this amount of propellant is vastly more than is required to do a nominal plane-change in LEO, the preferred approach is to use a moderate amount of propellant to change the plane of the orbit of the dry hardware from the Cape Canaveral orbit (28.5°, presumably) to the equator. At this point the full complement of propellant can be loaded on-board, and the vehicle can depart for Mars. This approach saves propellant by a significant factor over loading all the needed fuel at either L1 or the Cape Canaveral orbit.

Q: Are there ways to reduce the cost even further? A: It seems possible that mass-produced two-stage rockets burning LOX and liquid hydrogen could ultimately be launched from a tethered balloon at even lower cost than three-stage solid rockets, since liquid H_2/O_2 costs well under \$1/kg and has substantially higher performance. Perhaps the lowest-cost system will be a "blow-down" architecture, with boiloff pressure feeding propellants to the combustion chamber, avoiding the need for expensive turbopumps. Perhaps the (non-thrust-vector-controlled) combustion chamber and nozzle will be submerged directly in the base of the LH2 and LOX cryogenic tanks (respectively) to provide adequate pressure by sufficiently rapid boiloff of the cryogens, and to eliminate the expensive plumbing that traditionally circulates cryogens through fine tubes surrounding the nozzle and combustion chamber to cool them. The cryogens could be loaded into the inexpensive rocket shell from well-insulated tanks on the launcher as it makes the ~90-minute trip up to the balloon.

Q: Wouldn't it be better to make a high-flying aircraft to drop the rockets, instead of the tethered balloon? A: No. The development cost of a large "U-2 like" aircraft that is capable of dropping a 10-ton payload from an altitude of 26 km will be vastly greater than the cost of the tethered balloon system, and would transform this concept into "yet another" questionable idea that uses a large up-front capital investment to give a prospect of reducing operational cost. The benefit derived from the subsonic forward velocity of any economical aircraft is not substantial. A large investment (hundreds of millions of dollars) would have to be made in the aircraft before any substantial risks are retired. With the tethered balloon approach, a low-cost 1-cm diameter tether ($50K in materials cost) can be used to tether a commercially-available balloon on an equatorial island for an extended period to assess the seasonal wind profile and other environmental factors, retiring most system risks with very little investment. The energy and time required for an aircraft to climb to the proper altitude is substantial, and is a major negative in the overall economics of such an approach. Development of a spherical balloon that is sufficiently large and strong is relatively low risk given that comparable balloons are routinely flown by NASA, although later a blimp-shaped balloon that is higher cost and higher risk may be desired. Development of the payload module and head-end control system is moderate risk, but it pales in comparison to the complexity of subsystems developed for luxury automobiles that have acceptable cost when manufactured in similar volume. Development of moderate-performance (e.g. moderate propellant mass fraction) solid rocket stages is low risk. Once a prototype system is demonstrated with approximately the desired operational parameters, then an RFP can be let to industry for mass production. Industry can recover the cost of the factory and tooling by amortizing it over a volume production run in a manner similar to that for luxury automobiles. Probably the government would need to contract for a few years of guaranteed production to enjoy the benefits of this approach.

REFERENCES

1. A good discussion of the need for vast amounts of propellant in LEO to support the exploration vision is contained in the plenary talk "CA-1: Exploration Super System (ESS)" by Ralph Ewig of Andrews Space at the NASA Exploration Systems Technology & Systems Integration Technical Interchange Meeting, 13-17 December 2004, Ritz-Carlton, Pasadena, CA, CD proceedings disk1:Tuesday Plenary/6_Andrews Space CE&R.ppt (available from Neville Marzwell of the Jet Propulsion Laboratory).
2. See http://www.spacefuture.com/archive/designing_reusable_launch_vehicles_for_future_space_markets.shtml
3. See http://www.firefox-fx.com for on-line pricing and ordering information for rocket fuel ingredients including ammonium perchlorate and powdered aluminum.
4. See http://www.performancecomposites.com/design_guide2.html
5. See, for example, low-cost, GPS-augmented IMUs at http://www.xbow.com/
6. See http://www.solar.ifa.hawaii.edu/Tropical/summary.html
7. See http://www.physicalgeography.net/fundamentals/7q.html
8. See http://tdserver1.fnal.gov/nicol/lhc_irq_cryostat/ch_darve/public/publi/ICEC19_MLI.pdf
9. See http://www.digitaldutch.com/atmoscalc/
10. See http://www.yalecordage.com/html/industrial/double_ braids.html) for performance and pricing of commercial PBO rope.
11. Personal Communication, Jeff Hall (Jeffery.L.Hall@jpl.nasa.gov) and Andre Yavrouian (Andre.H.Yavrouian@jpl.nasa.gov) of the Jet Propulsion Laboratory.
12. See http://www.nsbf.nasa.gov/balloons.html
13. Presented at the IEEE International Conference on Robotics and Automation, Workshop on Integration of AI with Robotics, May 1989.
14. See http://www.nrl.navy.mil/content.php?P=ROCKOON
15. Joel Powell, "The China Lake Launches", *Smithsonian Air and Space*, Vol. 11, No 6, pp. 16-17, Feb/Mar 1997.
16. Peter Pesavento, "U.S. Navy's Untold Story of Space-Related Firsts", *Spaceflight*, Vol. 38, pp. 239-245, July 1996.
17. Tony Reichhardt, "The One-Pound Problem", *Smithsonian Air and Space*, Vol. 14, No 4, pp 50-57, Nov 1999.
18. Brian H. Wilcox, "Spin Control", *Smithsonian Air and Space*, Vol. 14, No 5, p 6, Dec 1999/Jan 2000.
19. Howard A. Wilcox, "Microminiaturized payloads open paths for space research", *Astronautics*, p 20, 95, May 1959.
20. See http://www.astronautix.com/lvs/propilot.htm
21. See http://www.tcomlp.com/overview_publications.html