Spaceport 2.0:

Cutting Launch Costs with Orbital Infrastructure

by

Roger Arnold

Principal Investigator for
Space Infrastructure

The Atlantis Project

Abstract

A concept for dramatically reducing the cost of launching passengers and cargo to space is introduced. It involves a spaceport orbiting in low Earth orbit (LEO), which we refer to as *Spaceport 2.0*[1]. Like the airports that are its terrestrial counterparts, it's a hub for scheduled arrivals and departures of flights. In this case, the arrivals are shuttle vehicles inbound from launch sites on Earth or from locations in cislunar space. The departures are flights returning to Earth or outbound to destinations in cislunar space.

The "runways" that Spaceport 2.0 employs to launch departing vehicles are linear electric catapults, similar in concept to the electromagnetic catapults that new aircraft carriers are using to launch fighter jets from their decks. They're much longer however. The launch catapults of Spaceport 2.0 are able to launch departing vehicles at speeds upwards of 3 kilometers per second. The runways employed to catch arriving vehicles are similar, but operate in reverse. In either case, the ultimate function of the catchers and launchers is the same: to provide ΔV to the shuttle vehicles with no expenditure of shuttle rocket propellant. That ends up making a large difference in the cost and practical feasibility of space transportation.

The concept introduced here allows affordable, high-capacity access to low Earth orbit and beyond. The infrastructure provided unlocks the door to a host of scientific and commercial applications that were once relegated to the pages of science fiction. This paper explains the concept and how it can be used to achieve the goals of Artemis and follow-on projects.

Part 1: Concept

A small conundrum

The depth of Earth's gravitational energy well is 62.56 megajoules per kilogram. That's 17.38 kilowatt hours, or $1.39 worth of electricity at an industrial electricity price of 8¢ / kWh. That's all it would cost to lift a kilo of mass from Earth's gravity well, *if* we had a way to do it using industrially priced electricity at 100% efficiency.

---

[1] "2.0" signifies that it's a modern update of a concept initially published in the late 1970s. See ref [1].

The cost of delivering a kilo of payload from Earth's gravity well using rockets is just a bit more than that. Around six orders of magnitude more, actually. That's owing to "the tyranny of the rocket equation" – the exponential increase in a rocket's mass ratio as its mission ΔV goes up in relation to exhaust velocity. The plots in Figure 1 tell the story.

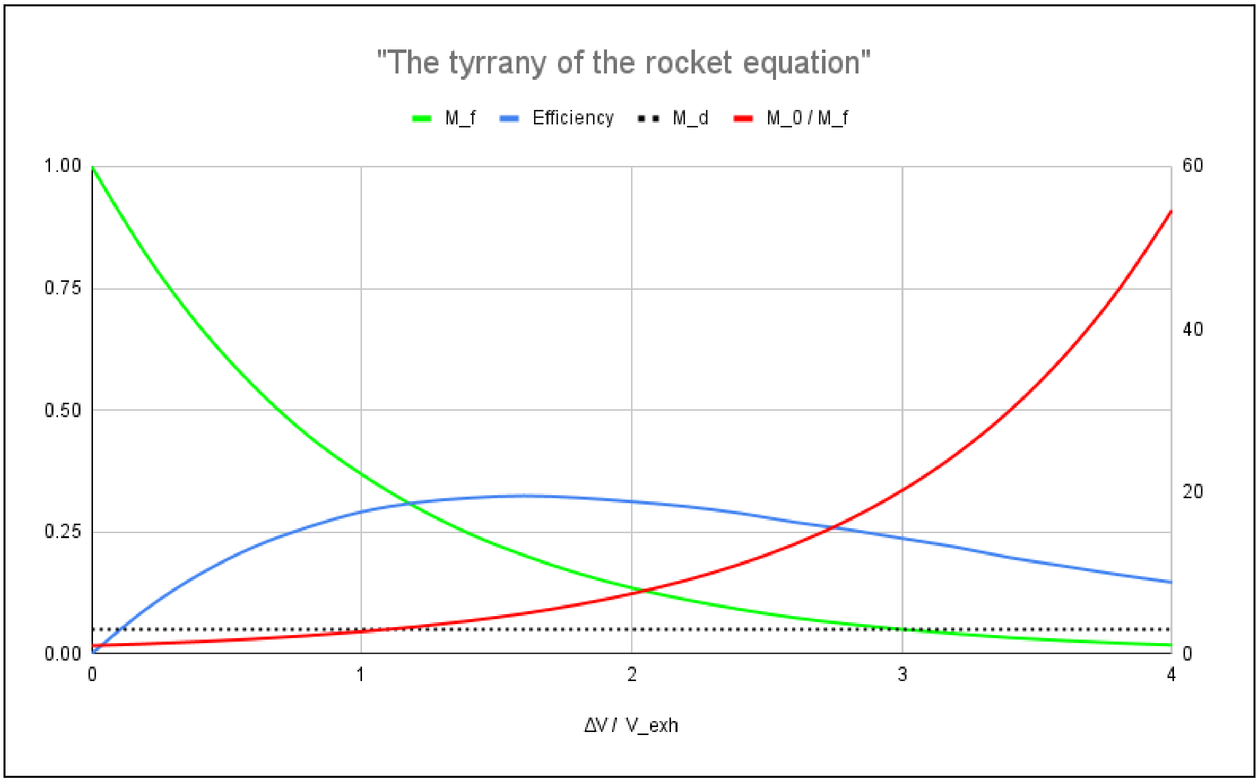

Figure 1

The horizontal axis for all of the lines plotted is the ratio of mission ΔV to exhaust velocity. The red line is the *mass ratio* for the rocket, conventionally m_0/m_f in the Tsiolkovsky rocket equation. The m_0 term is the mass of the rocket at T_0 (i.e., launch), while m_f is the final mass at engine shutdown. It's the rocket's dry mass plus payload. The dry mass is typically at least 5% of the wet (launch) mass. The chart shows that at that ratio, the maximum ΔV that can be achieved is three times the exhaust velocity. At that point the payload will be zero. To deliver any payload at a higher velocity, multiple stages must be used.

The dashed horizontal line in Figure 1 is the rocket's assumed dry mass. It's set in this case at 5% of m_0. The distance between that line and the green line of inverse mass ratio (i.e., m_f, taking m_0 as one) is payload. Where the green line crosses the dashed black line, the payload has fallen to zero.

The blue line is interesting. It's the fraction of the chemical potential energy in the rocket fuel that ends up as kinetic energy in the rocket's dry mass and payload. It's not widely known, but there's a region of mission ΔV in which rockets are reasonably efficient in converting chemical potential energy of fuel into kinetic energy of rocket and payload. Unfortunately that region of efficiency is limited to an interval from about 1 up to about 2.5 times the engine exhaust velocity. That range is adequate for suborbital flights beyond the Karman line, but not for orbital flights to LEO – much less for missions beyond LEO. That doesn't mean that chemical rockets can't be used to reach LEO and beyond; obviously they can. It just means they must use multiple stages with exponentially declining overall efficiency. Alas, chemical rockets are the only game in town, if we want to play in space.

Or are they?

Science fiction has anti-gravity and reactionless space drives. Unfortunately, they remain firmly in the realm of science fiction. We've no idea how to produce anti-gravity, and claims for various supposed reactionless drives have not proved out.

Be that as it may, scientifically credible alternatives to chemical rockets *do* exist. Included are space elevators, rotating tethers, and nuclear rockets. Each of those has its advocates, and there's much that could be written about them. Readers will have no trouble finding relevant information on the internet. For space elevators, in particular, there's an excellent website maintained by ISEC, the International Space Elevator Consortium.

In this author's opinion, none of the above alternatives have any realistic hope to be deployed within a time frame that would advance the objectives of NASA's Artemis program. In a competitive race to establish a permanent base on the moon, they would become operational too late, or with throughput capacities too limited, to do the job. There is, however, another alternative that might be deployed soon enough, and with sufficient throughput, to serve.

Hello space catapults

Linear electric motors (or catapults) are not 100% efficient in converting electrical to kinetic energy, but they're not far from it. If we can use them to deliver at least part of the energy we need to break the bonds of Earth's gravity, it will move the needle a long way toward lower launch costs.

If Earth were a smooth billiard ball with no atmosphere, we could build an electric catapult 2000 miles long (or half of Earth's radius) delivering a steady horizontal acceleration of 1.0 G. The catapult's payload would exit its muzzle at low Earth orbital velocity (7.8 km/s). Double the acceleration to 2.0 G, and the payload would exit at Earth escape velocity (~11 km/s). For any given velocity, there's a linear tradeoff between acceleration and length of the catapult. For robust payloads that can tolerate 1000 G's acceleration, a catapult would only need 2 miles to deliver orbital velocity. But it would be a space cannon, unusable for human passengers.

All that would seem academic in any case. Earth is *not* a giant billiard ball on which it's feasible to lay out arbitrarily long catapults. Also, Earth has an atmosphere. In the lower troposphere, the atmosphere is dense enough to reduce any payload moving through it at orbital velocity to a bright streak of plasma in a matter of seconds. Hence, unless we find a way to build the catapult *above* the atmosphere, acceleration to orbital or escape velocity by electric catapult is out.

Note that conditional, “unless we find a way”. Let’s look at that. Might there actually be a way we could build a linear electric catapult *above* the atmosphere?

Looking outside the box, LEO is above the atmosphere; could we build the required catapult there, in orbit?

The answer: sure, we *could* – in principle. But if the catapult is already *in* orbit, how on Earth (note, *on Earth*) can we use it to help us *get* to orbit?

A little more out-of-the-box seeking yields an answer. We can launch the shuttle on a suborbital trajectory whose apex is tangent to the orbit of the catapult. It’s a lot easier for a shuttle rocket to reach orbital *altitude* on a suborbital trajectory than it is to reach actual orbit. If we can get the timing right, our shuttle will touch the catapult’s orbit just as the catapult’s forward end passes. The shuttle gets caught by a “shuttle catcher” sled that’s been accelerated to match speed with it. Following the catch, shuttle and sled are decelerated together (relative to the catapult). When they come to rest near the aft end of the catapult, the shuttle has been successfully hauled into orbit, using momentum “borrowed” from the much heavier spaceport.

On arrival at the aft station, passengers disembark and cargo is offloaded. The empty shuttle is moved to the forward station, where passengers and cargo bound for Earth are loaded. After loading is complete, the shuttle is secured into its launch sled. Sled and shuttle then accelerate along the return-to-Earth track of the catapult. Just before reaching the end of the track, the sled hits the braking section of the track and the shuttle flies free. The sled will be returned to the forward station to launch the next shuttle back to Earth. The shuttle it has just launched will fall Earthward, continuing the suborbital trajectory that its capture by the spaceport interrupted. In the process, some or all of the momentum borrowed from the spaceport when the shuttle was hauled into orbit gets repaid.

The electromagnetic accelerator and braking tracks for catching arriving shuttles and launching departing shuttles are part of the spaceport’s *catcher-launcher* (C-L) system. They run along the chords of a long *C-L truss* that comprises most of the mass of the spaceport. But the tracks for catching shuttles from Earth and launching them for their return are only half of the C-L system.

Onward and upward

The other part of the C-L system is for transferring cargo and passengers between Spaceport 2.0 and destinations in cislunar space. From the aft station of the spaceport, launch and catch tracks for high orbit run to the forward end. The spaceport catapult is sufficiently long and powerful to put a lunar shuttle on a transfer orbit direct to the moon, starting from LEO. The diagram in Figure 2 below depicts the balancing momentum and energy flows.

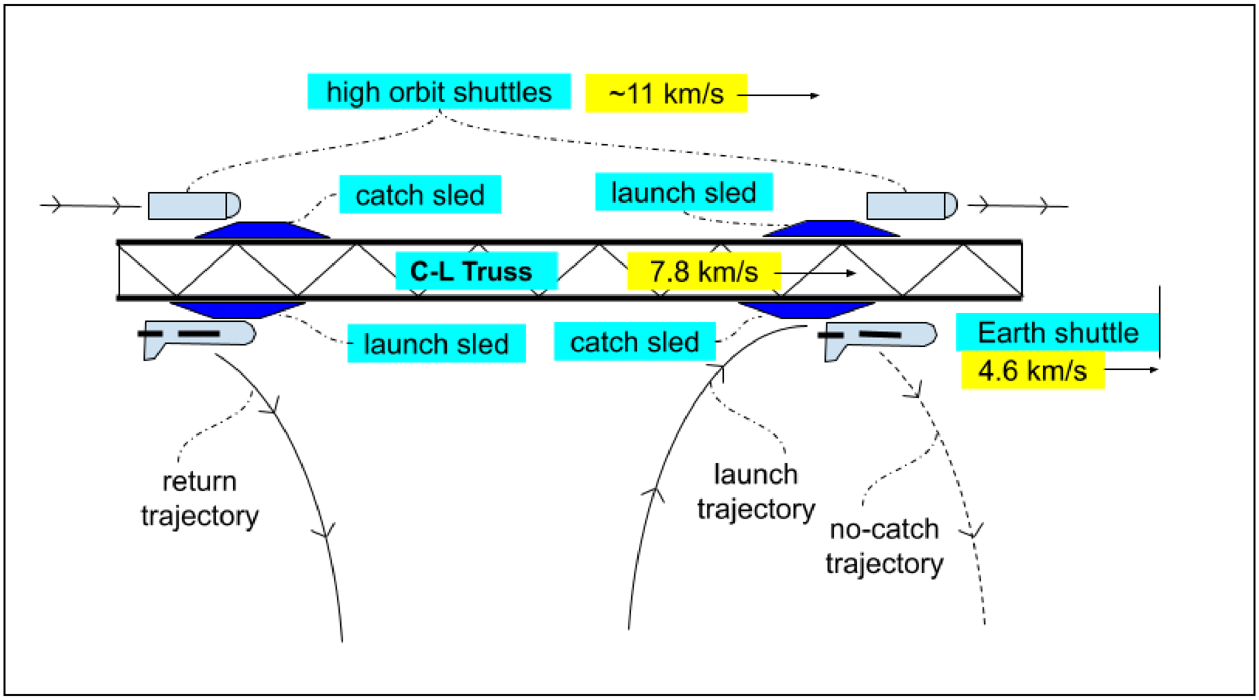

Figure 2: Catch and launch trajectories and velocities

A functional "replicator"

The lower loop depicted in Figure 2 reduces the rocket-supplied ΔV needed to reach LEO to the point that it can be supplied efficiently by reusable shuttles. It allows a 4:1 reduction in the propellant to payload ratio. However, it's in the *upper* loop where the really major payoff is realized.

The high orbit shuttles that are launched from the upper catapult rails can get 100% of the ΔV required for injection into a cislunar transfer orbit from the spaceport's C-L system. The only propellant the shuttles need is for control thrusters, plus a small amount for the insertion burn at their destination in cislunar space. If their destination is the lunar surface – and if a lunar C-L system has not yet been built – then they'll need fuel for landing and takeoff as well. But that will still be substantially less than they would need if they had to rely on rockets for injection to the cislunar transfer orbit. Similarly, shuttles returning from high orbit need not rely on high speed aerobraking and atmospheric reentry. They can deposit 100% of the ΔV they require to return to LEO in the catcher leg of the spaceport's C-L system.

The ΔV values to launch outbound high orbit shuttles and to catch inbound shuttles are the same. If outbound and inbound shuttles also carry the same mass, then there will be a net balance in both energy and momentum between inbound and outbound traffic. The spaceport's mass and the energy storage units of its C-L system provide short term buffering to smooth over the "lumpy" nature of arrivals and departures. But the outbound shuttles will mostly be carrying equipment and supplies for the bases; won't they mostly be returning empty? How can returning high orbit shuttles carry the same mass as outbound shuttles?

There's an easy answer. All that's needed is to load up returning shuttles with a sufficient mass of lunar rocks or regolith to match the mass the shuttles brought up. The loaded shuttles then return to the LEO spaceport, delivering the energy and momentum required to launch more outbound shuttles. The downward transport of lunar rocks and regolith balances the upward transport of equipment, supplies, and personnel.

That capability has profound implications for lunar bases and space habitats. It has long been assumed that the extreme cost of rocket-based transport of material to the moon would severely impede the development of lunar bases. Seeds of outposts might be planted at great cost, but growth would depend on the development of ISRU (In-Situ Resource Utilization) capabilities to make the outpost almost fully self-sufficient. That's a tall order. But with Spaceport 2.0 and its C-L system in the loop, the picture changes.

To the crew at the upper end of a high orbit shuttle loop, the loop serves as the functional equivalent of a Star Trek "replicator" device. In the Star Trek universe, the replicator was a universal recycler. Waste materials of any sort could be fed into the device. From the mass of the waste material fed into it, the replicator was able to fabricate an equivalent mass of anything whose pattern had been stored in the replicator's memory. The high orbit shuttle loop doesn't magically fabricate the items it delivers from the mass of "waste" material fed into it, but does the next best thing. It swaps the matter fed to it for an equivalent mass of items procured from the LEO spaceport.

The items procured could be things produced in space manufacturing facilities associated with the spaceport, or produced on Earth and delivered to the spaceport. The transport of material from Spaceport 2.0 to the cislunar receiving depot at the top of the transfer orbit is rendered essentially free, paid for by the energy and momentum of "waste" mass carried by returning shuttles.

Physical configuration

Figure 3 shows an end view of the C-L truss. It's a "truss of trusses", with six main chords (longitudinal stringers), which are themselves triangular trusses. As shown in the figure, the main chords are arranged in a hexagonal pattern, with two of the six planes defined by the main chords (the top and bottom) parallel to the surface of the Earth below. The maglev tracks for acceleration and deceleration of the catcher and launcher sleds run along the main chords. The two on top carry the catcher and launcher tracks for traffic from and to high orbit / cislunar space. The two on the bottom carry the catcher and launcher tracks for traffic from and to Earth. The two remaining main chords on either side of the truss carry utility tracks. Haulers running on those tracks carry passenger and cargo between the fore and aft stations of the spaceport. Energy storage Units (ESUs) are distributed along the length of the truss for buffering power to and from the launcher and catcher rails.

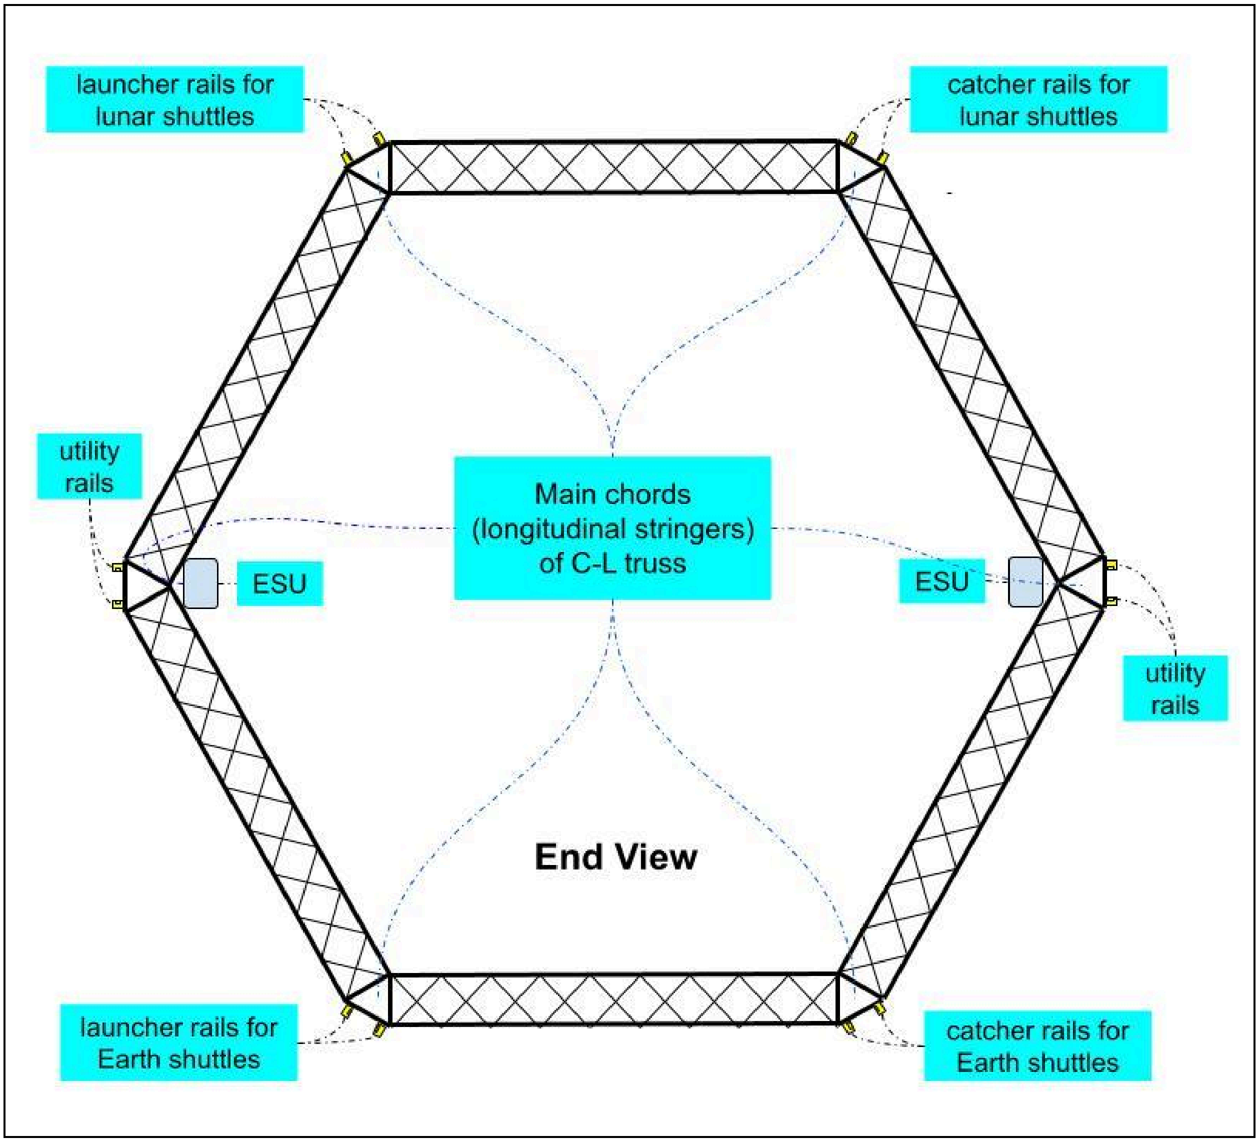

Figure 3: End view of C-L truss

Part 2: Feasibility

A key claim that we make for Spaceport 2.0 is that it can be built using basic technology that is available today. That's not to say that everything required is available off-the-shelf. An obvious case, for instance, would be the robots for in-space assembly of the C-L truss system. Their design and specific plans for how to use them will have to be developed. But there's no reason to think that they'll involve anything fundamentally beyond what can already be done with commercial robots. Neither the spaceport stations nor the C-L system require breakthrough developments in structural materials, power electronics, or high temperature superconductors. We could develop the design today and start building it tomorrow (figuratively speaking).

Building Spaceport 2.0 will be an ambitious undertaking. As a prelude for a future space-faring civilization, there will be many challenges. Some we can foresee now; others we'll encounter as work progresses. As is often the case, the challenges we can foresee will be the easiest to deal with; we'll work out viable solutions in advance. The ones that cause trouble are the ones we didn't anticipate. Those are the kind that lead to schedule slips and cost overruns. Project managers in the aerospace industry are familiar with that. There are standards and procedures to help mitigate the risk. Formal design reviews by peers and detailed design walk-throughs are often a pain for the engineers involved, but they pay off.

In that spirit, here's a short list of some foreseen challenges:

1) *Inherent instability*. The horizontal orientation for a long skinny object in orbit is unstable. The natural orientation is vertical, aligned with the gravity gradient. It requires active control to maintain the unstable equilibrium of the horizontal orientation.

2) *Reliability of capture*. The feasibility of safely and reliability "catching" a massive object with a closing velocity of kilometers per second are unproven. If the trajectory of an arriving shuttle is the least bit off, the shuttle might collide with the spaceport. The result would be catastrophic.

3) *Construction period*: The scale and complexity of the project to build Spaceport 2.0 suggests that it could easily take a decade or more to build. Throughout that period, it would be consuming a budget and resources comparable to what we put into the Apollo program in the 1960's. Can public interest and the necessary level of public funding be sustained over the course of a decade or more and multiple administrations?

Let's briefly examine each of these challenges in turn.

Active control

It's true that active control is needed to maintain the unstable equilibrium of the horizontal orientation of the C-L truss in orbit. However, as aerospace systems go, the dynamic control needed here is not a difficult thing to implement. The critical response time for stabilization of the C-L truss is orders of magnitude larger than it is for active stabilization of flight orientation routinely used in modern military aircraft. There are a number of avenues available to address stabilization, so fallbacks and failsafe modes are not difficult to provide for.

To maintain a horizontal orientation, the truss as a whole must rotate at the same rate at which it orbits the Earth. I.e, exactly one rotation per orbit. The easiest way to control an object's rate of rotation is to control its moment of rotational inertia. That's what figure skaters are doing when they pull in their arms to speed up a spin. The C-L truss will be an active structure. It will be able to adjust the tension and compression in its structural elements. It will be able to marginally extend or contract opposing chords to produce controlled bending and to adjust its moment of rotational inertia. With horizon sensors to detect small deviations of the truss from horizontal, fine control of attitude in orbit can be achieved with no expenditure of propellant.

There are many alternatives and backups that can be used for attitude control of structures in space. They're well known to aerospace engineers, and we needn't delve into them here. One worth mentioning, however, is the use of high impulse ion thruster engines. The C-L truss is equipped with thousands of such thrusters distributed along its length. Their primary purpose is fine control of the spaceport's orbit in the face of unbalanced captures and launches; they can also serve to provide a redundant source of attitude control for the spaceport and C-L truss as a whole.

Safety and reliability of capture

First, a point of clarification. When we talk about "reliability of capture" we're not talking about the percentage of shuttle launches from Earth that end up successfully docked at the receiving station of the orbiting spaceport. We expect that percentage to be high, but it need not be 100%. What's required is that the catcher system be 100% failsafe. Failed catch attempts are OK, so long as the manner in which they fail is safe.

The reason that the idea of catching a heavy shuttle vehicle that's approaching with a closing velocity of kilometers per second is so scary is that we're accustomed to thinking in terms of human senses and reflexes. But with modulated laser beams, interferometry, and precision time references, it's possible to design a monitoring system that tracks the shuttle position to within a millimeter at any nanosecond of time. The shuttle's trajectory is monitored from the moment of lift off. Any deviation from the expected trajectory is noted. Control signals are sent, and If the shuttle fails to respond properly, it indicates a problem somewhere. The catch attempt is aborted, and the shuttle returns to its launch site.

At no point is the shuttle allowed to be on a trajectory that will put it on a collision course with the space station. By main engine cutoff, the shuttle will be on a trajectory that falls just short of the catch trajectory. The last few meters per second of Delta V to a perfect catch trajectory are delivered by control thrusters, commanded by the spaceport's shuttle monitoring system. The shuttle aims for rendezvous, not with the catcher sled directly, but with a virtual twin of itself on a computed perfect trajectory. That computation is done by integrating the equations of motion backward from the instant of capture by the sled. When the shuttle has matched position with its virtual twin in state space, its thrusters shut off. It may still be a hundred kilometers and tens of seconds from the actual catch, but it has reached the exact point in the 12-dimensional state space of position and velocity from which it will coast to a smooth catch. With all thrusters off, no last-minute thruster failure can derail a smooth catch.

Of course, something could still go wrong with the catcher sled that would cause it to miss the catch. But that's not a catastrophe. The catch is programmed to occur at the apex of the shuttle's suborbital trajectory. If the catcher sled isn't in the right position with the right velocity to effect the catch, it won't happen. But the point where the catch would have occurred was the apex of a suborbital trajectory. The uncaptured shuttle will just fall back toward Earth on a continuation of its suborbital trajectory.

Construction period

An important consideration is that Spaceport 2.0 and its C-L truss can be developed in stages. It isn't necessary for the full 100 km orbital system to be completed for the system to be usable. An initial version with a length of 1.0 km would reduce a crewed shuttle's ΔV to LEO by only about 300 mps, but that's enough to increase the shuttle's payload to LEO by more than 10%. (See Figure 4).

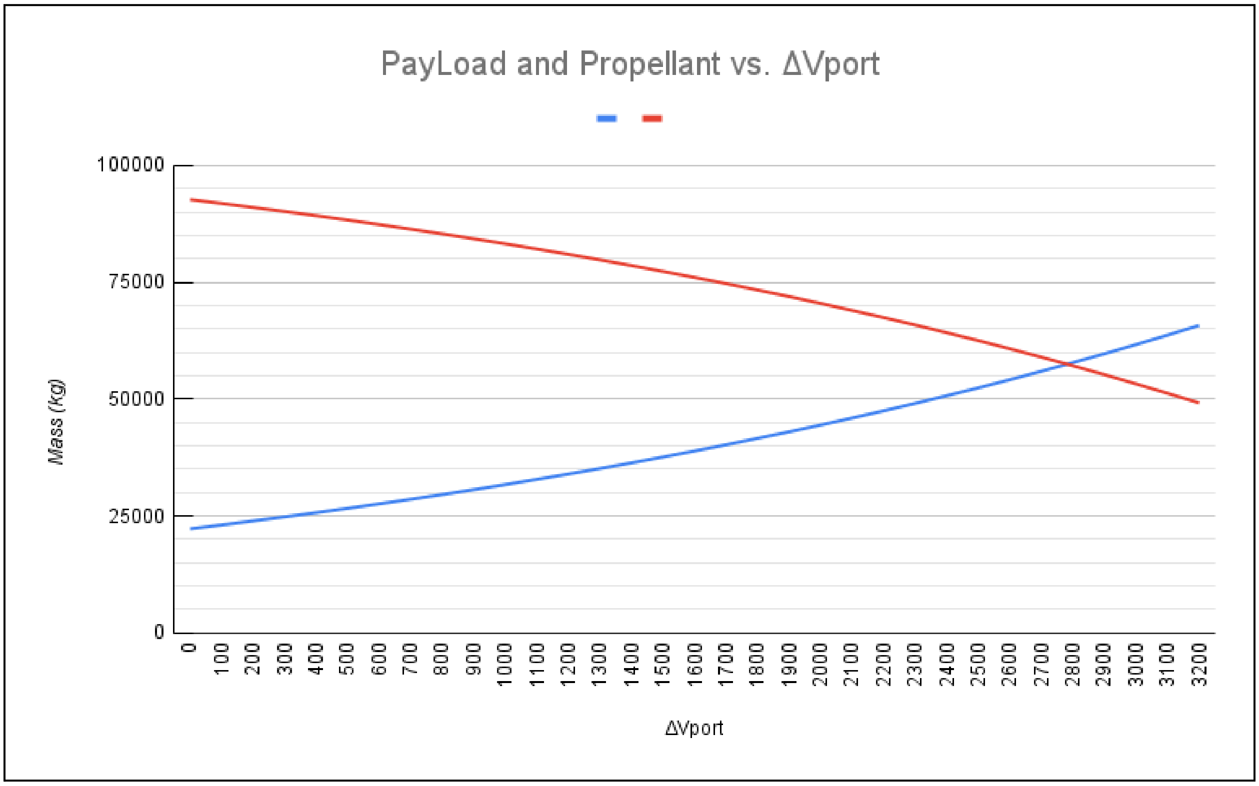

Figure 4

For uncrewed cargo deliveries, the benefits of an early, abbreviated version of the C-L system can be much greater than a 10% increase in payload. For a payload package that can tolerate 400 G's, a 1.0 km truss could deliver the full 3200 mps from the apex of the suborbital trajectory to full orbital velocity. Likewise, the abbreviated truss could deliver the full 3200 mps required to inject the payload into a cislunar transfer orbit. In both cases, the mass of the high acceleration shuttle plus payload would be limited to 1% of the final design target for the completed C-L truss. However if the final design target is 100 tons, the capability to deliver a ton of payload to LEO and then on to cislunar space for the ΔV of a modest suborbital flight would be quite useful.

The upshot is that with proper planning for it, Spaceport 2.0 can begin paying dividends in lowering the delivery cost of equipment and supplies to the Lunar Gateway Station – or any other destination in cislunar space – from an early point in its development. The cost reduction and high throughput possible will be game-changing to the scope of the lunar base that it's feasible to deploy.

Conclusion

Spaceport 2.0 effectively moves cislunar space “next door” to LEO (in terms of rocket Delta V), while at the same time moving LEO much closer to Earth (in the same terms of delta V). The result is a 40:1 reduction in the rocket propellant required to deliver a ton of payload to any destination in cislunar space. Without the capabilities that Spaceport 2.0 brings, talk of a permanently staffed lunar base on a scale comparable to, say, the Amundsen-Scott South Pole Station, is only talk. Such a base would be far too expensive to set up and supply by chemical rockets. And what’s true for a lunar base would be doubly so for a Martian outpost. Spaceport 2.0 opens the door for those possibilities and more. Deep space habitats and asteroid mining suddenly become serious possibilities.

Spaceport 2.0 *can* be built using basic technology that is available today. Its eventual goal, as envisioned here, is the ability to impart 3200 mps of ΔV electromagnetically at a crew-friendly 4 G’s of acceleration. That will enable crewed shuttles to be launched to cislunar injection orbits from LEO with no expenditure of rocket fuel. Reaching that level of development may be the work of a decade, but the system allows for cheap cargo delivery and bootstrapping well before that ultimate level of development is reached. In only a few years from commencement, it would be capable of efficiently provisioning a lunar gateway station or permanent lunar bases. That’s how a practical journey on the path toward a multi-planetary civilization can begin.