Is Lumen Orbit (Now Starcloud) Cooked? Order of Magnitude Estimates for Training AI in Space

Marcus Tsuei

Published Dec 2024, reviewed by Dr. Igor Bargatin

Introduction:

In recent years, advancements in the field of artificial intelligence (AI) have caused a surge in demand for the computing power needed to train and develop more advanced AI models. Data centers filled with racks of graphics processing units (GPUs) process terabytes of data which train and enhance the performance of AI models.

Data centers consume large amounts of electricity and power consumption for data centers in the US is expected to increase ~800% to 652 TWh by 2030.[1] For comparison, the demand would make up roughly 16% of current U.S. electricity demand. There is now growing concern from the big tech companies that develop and train AI models over whether or not the energy infrastructure can even support the projected demand. Microsoft is going to lengths such as reopening the Three Mile Island nuclear plant, the source of the U.S.'s worst commercial nuclear plant disaster, in order to power its data centers.[2]

A proposed solution to providing enough energy is to move data centers to space in order to take advantage of abundant solar energy and passive radiative cooling. A startup called Lumen Orbit is proposing to deploy a large gigawatt scale data center in orbit. According to their calculations, they would deploy these data centers in several shipping container styled spacecraft, along with a 4km x 4km solar panel array.[3]

Building the system proposed by Lumen Orbit will likely be the largest space structure ever built. This paper will validate the assumptions made by Lumen Orbit and factor in additional considerations in order to determine if an orbital data center is a reasonable solution for training AI models in the future.

Required Solar Power for each GPU rack

In orbit, all of the electricity supplied to the datacenter will likely come from solar panels. Since GPUs are packaged into racks, we can model the compute capability of a data center by its number of racks and scale power consumption accordingly. The newest GPU rack from NVIDIA, the GB200 NVL72 operates at 120kW of power.[4] The solar irradiance above the atmosphere, or energy density of the sun's rays, is 1.36kW per square meter.[5] Solar panels are not perfectly efficient at converting solar energy to electricity, and solar panels used in modern satellites have an efficiency of about 32%.[6]

Thus, the total surface area needed can be expressed as

$$\frac{Required\ Power}{\eta_{solar} \cdot Irradiance} = \frac{120 \times 10^{3}\ kW}{(0.32) \cdot (1.36 \times 10^{3} \frac{W}{m^{2}})} = 275.74\ m^{2}/rack$$

---

[1] https://www.forbes.com/sites/bethkindig/2024/06/20/ai-power-consumption-rapidly-becoming-mission-critical/
[2] https://www.npr.org/2024/09/20/nx-s1-5120581/three-mile-island-nuclear-power-plant-microsoft-ai
[3] https://lumenorbit.github.io/wp.pdf
[4] https://www.nvidia.com/en-us/data-center/gb200-nvl2/
[5] https://en.wikipedia.org/wiki/Solar_constant#cite_note-KoppLean11-1
[6]https://www.nasa.gov/smallsat-institute/sst-soa/power-subsystems/#:~:text=However%2C%20in%20the%20aerospace%20industry,on%20the%2 0solar%20cells%20chosen.

Discrepancy in ISS solar panel array size and theoretical array size

For comparison, the International Space Station's (ISS) solar panels span 2500 m^2 and generate a similar 120kW.[7] It is peculiar that the ISS utilizes a solar array that's almost an order of magnitude larger than what's theoretically needed. There are a few explanations for this:

Firstly, the ISS relies on legacy solar panel technology that has an efficiency of only 14%.[8] Secondly, the ISS's orbit experiences darkness during half of its orbital period, reducing solar power by another 50%.[9]

Given a solar irradiance of 1360 Watts/m^2, 14% solar efficiency, efficiency, as well

$$Irradiance = (1360 \cdot 0.14) \cdot 0.50 = 102.2\, watts/m^2$$

At NASA, common practice is to oversize a solar array by 20% as a safety factor for sufficient power generation.[10] Factoring in these extra considerations, the required solar array for the ISS can be modeled as

$$\frac{Required\ Power}{\eta_{solar} \cdot Irradiance} \times\ Factor\ of\ Safety = \frac{120 \times 10^3\ kW}{(0.14 \cdot 0.50) \cdot (1.36 \times 10^3 \frac{W}{m^2})} \times\ 1.2\ = 1512.6\, m^2$$

Which falls within an order of magnitude of what the ISS currently uses. The ISS recently introduced a new solar array called the Roll Out Solar Array (ROSA) which packs solar cells more efficiently in the array and an increased solar conversion efficiency of 30%.[11]

The 6 ROSAs measure 82.2 m^2 and can source 20kW each which totals to 493 m^2 for 120kW and is closer to what is theoretically modeled.[12]

Understanding that shading will greatly reduce the efficacy of solar panels, it's imperative to find an orbit that can provide constant sunlight to minimize the additional mass required for extra solar panels and batteries to store the energy.

Radiative Heat Dissipation

Along with more solar energy, putting a data center in Space is an attractive option due its ability to act as a giant heat sink. In addition to powering a GPU rack itself, a data center also needs to provide sufficient cooling for the rack as almost all of the energy consumed by data centers is dissipated as heat.

---

[7] https://www.nasa.gov/image-article/solar-arrays-international-space-station-2/
[8] https://ntrs.nasa.gov/api/citations/20190032191/downloads/20190032191.pdf
[9] https://www.nasa.gov/wp-content/uploads/2017/05/569954main_astronaut20_faq.pdf
[10] https://ntrs.nasa.gov/api/citations/20180007969/downloads/20180007969.pdf
[11] https://ntrs.nasa.gov/api/citations/20190032191/downloads/20190032191.pdf
[12] https://ntrs.nasa.gov/api/citations/20190032191/downloads/20190032191.pdf

Since the atmosphere is so thin in orbit, the orbital data center will have to transfer the waste heat via radiation. Heat transfer by radiation is expressed as

$$P = \varepsilon\sigma A \Delta T^{4}$$

The required area is expressed as

$$A = \frac{P}{\varepsilon\sigma(T_1^4 - T_2^4)}$$

Where ε is the emissivity constant of the material that's radiating heat, σ is the Boltzmann constant, and ΔT is the temperature differential of the emitting surface and the surrounding environment.

Using values from the radiators on the ISS, the ε for the special Goddard Space Flight Center (GFSC) paint on the radiators is .92, σ is 5.67×10^{-8}.[13]

Although the heat transferred is radiated to the vacuum of Space, the heat would not be radiated directly from the GPU rack due to limited surface area of the chips.

If we were to radiate directly from the GPU and coated it in the GSFC paint, the required radiating area would be:

$$\frac{120\times10^{3}}{(.92)(5.67\times10^{-8})(310^{4}-2.7^{4})} = 249\ m^{2}/rack$$

Which is much larger than the entire surface area of all the components in a GPU rack [find a source for this. As a result, we would still use a liquid cooling loop. The ISS uses ammonia for its radiators and cools the ammonia to 233.15K.[14]

Thus, the required surface for the radiators will be:

$$\frac{120\times10^{3}}{(.92)(5.67\times10^{-8})(310^{4}-233.15^{4})} = 366.29\ m^{2}/\text{rack}$$

Ammonia is extremely toxic. While it is a favored chemical for cooling because of its low freezing temperature, scaling operations pose complications due to its hazardous nature. Manufacturing large radiators with ammonia loops will require a specialized containment facility comparable to the way SpaceX processes hypergolic rocket propellent (i.e. in the middle of nowhere in an old Air Force Bunker)

Additional Pump Energy Costs

Adding a liquid cooling loop means that the orbital data center will need to also power pumps. Although we still have to use a liquid cooling loop, operating in space is cold enough to allow for a single phase liquid ammonia cooling system.

To estimate the power needed to run the pumps[15]

$$P_{pump} = \frac{\Delta P \cdot \dot{V}}{\eta}$$

---

[13] https://ntrs.nasa.gov/api/citations/19840015630/downloads/19840015630.pdf
[14] https://www.nasa.gov/wp-content/uploads/2021/02/473486main_iss_atcs_overview.pdf
[15] https://neutrium.net/equipment/pump-power-calculation

According to the Collins Aerospace Datasheet, ΔP is 7.2 PSID or 46942 Pa.[16] In general η is around 85% for higher end centrifugal pumps.[17] For the volume flow rate, we can calculate the necessary mass flow to dissipate 120kW:

$$Q = \dot{m} \cdot c_p \cdot \Delta T$$

Where the specific heat capacity of liquid ammonia, c_p, is 4744 $J/kg \cdot K$, and the temperature differential, 77k, is modeled from the ISS radiators.[18] Solving for the mass flow

$$\dot{m} = \frac{Q}{c_p \cdot \Delta T} = \frac{120\times10^3}{4744 \cdot 77} = 0.328\ kg/s$$

And dividing by the density of liquid ammonia, 696 kg/m^3, the pump power needed to transfer heat to the radiators is

$$\frac{\Delta P \cdot \dot{V}}{\eta} = \frac{46942 \cdot (\frac{0.382}{696})}{0.85} = 30.31W$$

Estimating the cost to cool a ground based data center

Terrestrial data centers do not have the advantage of being able to radiate heat effectively because of a low temperature differential. For example, even in the winter, a data center in Arizona would get a temperature differential of 10 K. The data centers will need a large amount of surface area to radiate heat effectively.

As a result, most data centers will use chillers to reject the heat. The chillers utilize a refrigerant cycle and require power for compressors and other internal components. The energy required for chillers is measured by their coefficient of performance (COP).[19]

$$COP = \frac{Cooling\ Capacity}{Power\ Input}$$

The COP for a large commercial electric chiller is 6.5.[20] Thus for 120 kW of cooling, you would need to input 18.46 kW of energy to run the chillers. By this estimate, being able to radiate heat into space provides a significant advantage.

Mass to orbit for a 40 MW data center

Lumen Orbit's initial white paper hypothesizes launching a 40 MW data center into orbit. Using the previous calculations, the total payload mass can be estimated as

---

[16]https://www.collinsaerospace.com/-/media/CA/product-assets/marketing/a/active-thermal-control-systems-data-sheet.pdf
[17] https://neutrium.net/equipment/pump-power-calculation/
[18] https://www.engineeringtoolbox.com/ammonia-d_1413.html
[19] https://aircondlounge.com/cop-seer-eer-kw-ton-ceer-hvac-efficiency-guide
[20] ASHRAE via chatGPT, access is blocked by paywall

$$m_{total} = N \cdot m_{rack,\, total} + \ m_{additional} m_{solar} + m_{radiators} + m_{racks} + m_{container}$$

Where

$$m_{rack,\, total} = \ m_{solar} + m_{radiators} + m_{rack}$$

The mass of each NVIDIA GB200 rack is 1360 kg.[21]

The mass of the solar panel arrays can be derived from the mass of the ISS ROSA arrays. Each ROSA array has a mass of 325 kg, and a surface area of 82 m^2, resulting in a mass/surface area ratio of 3.96 kg/m^2.[22] If each rack requires 276 m^2 of solar panels, then 1092 kg of solar panels are needed for each rack.

The same approach can be applied for the radiators. A radiator panel is 79.22m^2 and has a mass of 1122 kg → mass/surface area ratio of 14.16 kg/m^2. Each rack requires 366.29 m^2 of radiators → 5186.6 kg per rack.

Each rack occupies a 1.433m^3 volume. An empty standard 40 ft shipping container weighs 3800 kg and can fit 66m^3 of cargo volume which equates to roughly 45 racks.[23] We can package our racks into a shipping container and use the weight of the shipping container to account for additional space hardware and a lighter enclosure structure. For reference, a medium sized satellite bus such as the Maxar 1300 weighs around 939 kg.[24]

Our final mass of the data center system per rack can be expressed as

$$m_{datacenter,total} = 7638 \cdot N \ + 3800 \cdot \frac{N}{45}\, kg/rack$$

Where N is the number of racks required.

Conclusion

To minimize shadowing of the solar array from the Earth, the orbital data center must follow a sun synchronous, dawn to dusk orbit. A dawn to dusk orbit follows the line that separates day and night so that the satellite will always be exposed to the sun.

In their white paper, Lumen Orbit proposes first launching a 40 MW data center. Using NVIDIA GB200s, a 40 MW data center would need ~ 333 120 kW racks and 8 shipping containers. Using the model of total mass to orbit, this data center would require

$$m_{datacenter,total} = 7638 \cdot 333 \ + 3800 \cdot \frac{333}{45} \frac{kg}{rack} = 2571 \times 10^{3} = 2571\, tons$$

[21] https://training.continuumlabs.ai/infrastructure/servers-and-chips/nvidia-gb200-nvl72
[22] https://www.nasa.gov/wp-content/uploads/2018/07/spacex_crs-11_mission_overview.pdf?emrc=bd7dce
[23]https://www.mobilemodularcontainers.com/blog/40-ft-container-dimensions#:~:text=Here's%20how%20you%20can%20calculate%20the%20CBM:,=%20Internal%20(Length%20x%20Width%20x%20Height)
[24] https://rsdo.gsfc.nasa.gov/images/catalog-rapidIV/Maxar_1300_Data_sheet-Rapid_IV.pdf

The ISS ROSA array upgrade cost $103 million for ~500$m^2$ of solar panels which breaks down to about $206,000/$m^2$ of solar. Let's assume that Lumen Orbit can bring down the cost of solar panel manufacturing by 1000x to $206 (which is close to what it costs for terrestrial solar), Total solar costs for the 40MW cluster would amount to 333 racks * 276 m^2 of solar panels per rack * $206 per m^2 = ~$18M

Lumen Orbit estimates that a 40 MW data center would only take 100 tons. Below is a cost analysis from Lumen Orbit's whitepaper that compares infrastructure needs.

Table 1. Cost comparison of a single 40 MW cluster operated for 10 years in space vs on land.

Cost Item	Terrestrial	Space
Energy (10 years)	$140m @ $0.04 per kWh	$2m cost of solar array
Launch	None	$5m (single launch of compute module, solar & radiators)
Cooling (chiller energy cost)	$7m @ 5% of overall power usage	More efficient cooling architecture taking advantage of higher ΔT in space
Water usage	1.7m tons @ 0.5L/kWh[10]	Not required
Enclosure (Satellite Bus/Building)	Approximately equivalent cost	
Backup power supply	$20m (commercial equipment pricing)	Not required
All other data center hardware	Approximately equivalent cost	
Radiation shielding	Not required	$1.2m @ 1kg of shielding per kW of compute and $30/kg launch cost
Cost Balance	**$167m**	**$8.2m**

An updated version of this table would look like

Cost Item	**Terrestrial**	**Space**
Energy	$140m	$18m (assuming Lumen can revolutionize solar panel manufacturing)
Launch	None	$125m
Cooling	$7m	Probably the same if you're going to manufacture enough radiators to cover 18 soccer fields

Water Usage	Hella water	Hella ammonia
Enclosure	same	same
Backup Power	$20m	none
All other data center hardware	same	same
Radiation shielding	none	$1.2m
Total Cost	$167m	$145.2m

Additional Considerations
Micrometeorite impacts could cause radiator leakage, Electrostatic Discharge Events that could cripple the whole solar array, Fuel for altitude adjustments, gravity gradient that occurs when you have a space station that's 4km long…

With some order of magnitude estimates, it seems like the biggest bottleneck still comes down to energy. To make this viable, Lumen Orbit will need to become the greatest solar panel and radiator manufacturer in the world.