SERVER SKY

Computation in Orbit

Keith Lofstrom keithl@ kl-ic.com

http://server-sky.com 2009 December 7

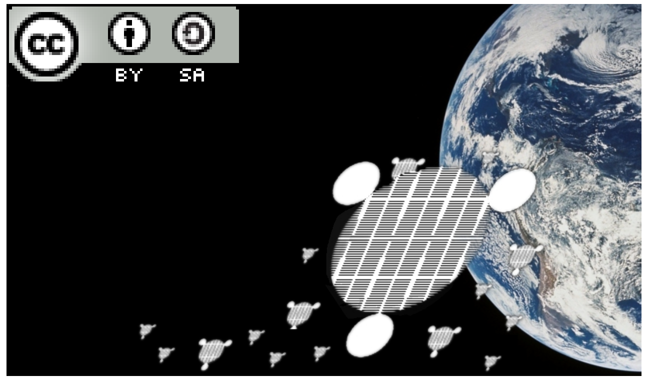

ABSTRACT

It is easier to move bits than atoms or energy. Server-sats are ultralight disks of silicon that convert sunlight into computation and communications. Powered by a large solar cell, propelled and steered by light pressure, networked and located by microwaves, and cooled by radiation into deep space. Arrays of tens of thousands of server-sats act as highly redundant computation and database servers, as well as phased array antennas to reach thousands of transceivers on the ground.

First generation server-sats are 20 centimeters across (about 8 inches), 0.1 millimeters (100 microns) thick, and weigh 7 grams. They can be mass produced with off-the-shelf semiconductor technologies. Gallium arsenide radio chips provide intra-array, inter-array, and ground communication, as well as precise location information. Server-sats are launched stacked by the thousands in solid cylinders, shrouded and vibration-isolated inside a traditional satellite bus.

A Story

You've been working in space for six years now. Of course, your body and home and family are on Lopez Island in Washington's Puget Sound, but you work with your mind. Your mind feels like it is thousands of kilometers away, in a sensor cloud in the Indian Ocean. Or in a huge cluster of tiny computing satellites currently passing into the Earth's shadow, somewhere over the Atlantic ocean east of Brazil. Or perhaps it is in the communication cluster over the mid-Pacific ocean directly south of you, 3000 kilometers west of Ecuador. That is where the house dish is pointing.

Now your mind comes home from work. It's time for your after-work bicycle ride. With most of your neighbors telecommuting, the island's narrow roads will be fairly clear, even at rush hour. They will stay clear until the robot delivery trucks come out tonight, doing their silent 15 mile-per-hour rounds. You put the violin you are sending your nephew in the mailbox. No package or address, you've told the service where it's going, and the robots and the web take care of the rest. You get on your bike.

You check YouNews on the head-mount net display - a pod of orcas are playing around Swift's Bay. You peddle over to Port Stanley. The identities of the orcas aren't established yet, but you expect to know who they are by the time you get there. An orca named George has been a frequent visitor to Lopez. Perhaps the web will even have the tracking numbers for the salmon George ate today. But not names. Naming individual salmon would be *silly*. Besides, with the hydropower dams removed, and all the planet's power coming from space, there are too many salmon to name anyway.

Orcas. You probably shouldn't be thinking about work; as a computational ecologist, you spend your working day in a model of the ecosystem off Somalia. You are looking for lingering radiation effects of the Mogadishu Nuke. The orcas at the top of the food chain in the Indian ocean seem to be concentrating strontium 90, and you are studying alternatives for delivering chelating agents to their food supply. But watching real live orcas playing in the water near home is a way to “decompute”, and get grounded on terra firma.

Some of your relatives spend 24 hours a day immersed in the computation sky, even during sleep, and some claim that the center of their consciousness has shifted there. It can seem that way, wielding virtual cognitive powers many times the power of your brain. But shifting into the sky? Nonsense. The emails you keep getting from your dead uncle Jack not withstanding. Yes, his roses are beautiful, and he is learning to create a meditation garden from a monk in Chang Mai. But his *parachute failed*, dammit. Jumping out of sub-orbitals at his age ...

You have access to billions of times more computing than was available to the whole world in 2010, but very little of it is in the homes and buildings around you. It is mostly in orbit around the Earth, with an increasing amount now building in the Jupiter orbit computing belt. The speed-of-light round trip is up to 3 hours, but they are running accurate one week weather simulations, so the delay is tolerable for the big jobs. Jack's wife Annie (also *dead*, dammit) wants to move there, but he wants to remain within a light-second of his garden. Virtual dirt just isn't the same, he says. Annie calls him old fashioned.

The desktop computer, laptops, PDAs, and other standalone computing devices have disappeared, replaced by sensors and communication links. The computers on your body are fast enough to measure you, and construct images for your senses, but mostly they are rendering scenes constructed thousands of kilometers away, selecting content branches based on your last second of activity. The personal-area computing power is high by 2010 standards, but it is finite, you can only see and feel so fast. With computation efficiency doubling every two years, the sensors just get smaller and less obtrusive.

Night is coming, time to ride home for supper, before the rainstorm scheduled at 7:30. You could ride with the head-mount active, and the roads would seem as bright as day. Instead, you'll enjoy the evening light. Some folks hate the glow from the lunar computation band. It washes out the Milky Way, they say. But it is not nearly as bright as the full moon, and with most outdoor lights gone, the sky is a lot darker than it was in your childhood. If you need to see a hazard in a hurry, your head-mount will come on and highlight it for you. The sky-glow is faint, but it lights the road home.

Solving the Computing Energy Crisis

Traditional data centers consume almost 3% of US electrical power, and this fraction doubles every five years [DATA]. Computer technology is improving - new hardware can deliver the same computation for half the power of two-year-old hardware. But the demand for computation is increasing more rapidly.

Most of the computing growth is occurring outside of the United States, in rapidly developing countries such as China. Some estimate the total computing power for the planet is doubling every year, implying that world computing energy demand doubles every two. We are not constructing nearly enough clean power plants to meet this rapidly growing demand, and the competition over a limited amount of fuel for them will become increasingly deadly in the coming decades. The U.S. may have less total generating capacity in 20 years than we do now, while data center and data communication power usage may have increased to 40% of total load - or more.

One likely outcome is power rationing. In the best case, virtualized computers will be given smaller and smaller time slices on crowded hosts, increasing response time. Fiber internet to the home is capable of enormous bandwidth, but the optical network terminals at the customer end and the switches and routers at the ISP end may need to be slowed down to reduce power, also increasing response time. Of course, the demand for digital semiconductors will flatten out, mostly energy efficient replacement parts.

The worst case is more likely to resemble the historical case. During the California energy crisis, utilities reacted to high demand by shutting down power to customers. While data centers are often powered through battery-backed uninterruptable power supplies, these systems are limited, expensive, and inefficient, and may result in data centers restricting the compute load during blackouts. More likely, the data centers will go dark after sufficiently long power outages.

Packets travel through dozens of switches between the data center and the end user. While the internet is agile, and can route around failed links, a sufficient number of failed switches may result in inefficient routes, increasing the burden on the switches that remain. The result will be an increasingly slow, unreliable, and unpredictable internet. With "smart power" grids becoming increasingly dependent on computing and internet communication to extract maximum efficiency from limited generation, we may get into deadly positive feedback loops, leading to cascading failure of the combined computing and generation grid.

Alternative energy systems such as ground-based solar photovoltaic have been proposed, but solar panels intercept sunlight that would otherwise feed the biosphere. Generating the world's energy needs (estimated at 40 Terawatts by 2050 [SMAL]) with limited efficiency solar cells will require millions square miles of solar arrays. Estimated roof area for the entire United States is about 30,000 square miles, and paved area is around 60,000 square miles [AREA]. Covering many times that area with solar collectors will be proportionally more expensive than all our roads and buildings. Probably much more expensive, because roads are not made of fragile solar cells and do not need to collect, transform, and transmit electricity. Most importantly, solar power goes away at night - storing 12 hours worth of electrical generation will also require huge amounts of infrastructure. Terrestrial solar energy is interesting, and useful for small and remote systems, but it is not a practical way to generate Terawatt levels of electricity.

Space is full of energy from the Sun. Space solar power satellites [SSPS] are proposed to capture some of this energy and beam it to earth. SSPS captures solar energy in space, turns it into intense microwave beams, and focuses them on large "rectennas" on the ground, where the energy is converted back to electricity for local use. If the satellites are in geosynchronous orbit, the beam-spread at the ground will be large, so the rectennas must be large as well. However, they can be built to be radio-opaque while passing most of the visible light, so the area underneath can be used for growing crops. This provides high quality electrical energy without too much waste heat dissipated or land lost in the biosphere, if the rectennas are designed efficiently. Some of this power can drive data centers.

However, there are many inefficiencies between the solar cells in orbit, and the actual computing circuitry in the data center. Here are some of the steps and some estimated (WAG) efficiencies:

System	component efficiency (WAG)	cumulative efficiency	cumulative power (w/m^2)
Solar light input			1300.0
Solar cell on satellite	0.20	0.200	260.0
Power bus on satellite	0.80	0.160	208.0
Conversion to microwaves	0.50	0.080	104.0
Beaming efficiency	0.90	0.072	93.6
Atmospheric transparency	0.90	0.065	84.2
Collection efficiency	0.90	0.058	75.8
Rectification to electricity	0.60	0.035	45.5
Transmission in array	0.90	0.031	40.9
Conversion to high voltage	0.80	0.025	32.8
Transmission lines and substations	0.80	0.020	26.2
Data center HVAC, UPS, conversion	0.50	0.010	13.1
PC internal conversion to 1V CPU power	0.70	0.070	9.2
Available to compute load			**9.2**

Of the 260 watts of electricity produced by the solar cell, about 3.5% reaches the compute load. The rest is lost along the way. This does not include the power used to orbit the solar power satellite, or to keep it supplied with orientation fuel, so overall system efficiency is lower.

What if the conversion steps between the solar cell and the compute load could be eliminated, and all 260 watts per square meter could be turned into computation? One way to do that is to move the computer and the data center functions into space, The solar cell directly produces the high current. low voltage power that a CPU needs. The cost-effectiveness of space-generated power goes up by almost 30 times.

If the data are radioed to and from points on the Earth, much of the power and resource-consuming communications infrastructure can be eliminated as well.

System	component efficiency (WAG)	cumulative efficiency	cumulative power (w/m^2)
Solar light input			1300.0
Solar cell on satellite, used centimeters away	0.20	0.200	260.0
Available to compute load			**260.0**

Server Sky is one way to compute directly with solar energy. It strips away the mechanical structure, power transmission and conversion, and large power transmitters of a solar power satellite, so it is much cheaper to launch and easy to make.

Server Sky

Server sky is multiple arrays of ultra-thin (100μm) 7 gram satellites, maneuvered by light pressure, each of which converts electricity from a 15cm (6 inch) solar cell directly into computation and radio transceiver power. Server satellites are mass produced by the millions or billions, and are launched in stacks of tens of thousands with conventional rockets into a 6411 kilometer altitude orbit. The satellites deploy into large arrays to form phased array radio beams that can address many small spots on the ground. Recent advances in distributed array computing, CMOS radiation resistance, error detection and re-computation, and electro-chromic light shutters allow server-sats to be cheaply manufacturable with existing factories, some idled by recent economic troubles. Although expensive to launch, they will quickly pay for the launch cost through power and infrastructure savings.

In the longer term, electrically-powered launch systems such as the Launch Loop or the Space Elevator can reduce launch cost by orders of magnitude. Server-sats can be reconfigured to beam energy to the ground, like solar power satellites, and that power can lower the cost of launching more server-sats. Within a few decades, Server Sky can replace most ground-based computation and power generation, providing the entire world with first-world energy and information access.

Server Satellites

Silicon circuits, solar cells, and interconnect are essentially two-dimensional systems. The horizontal dimensions of an integrated circuit die may be measured in millimeters, but all the important action occurs within a few microns of the top surface. Indeed, modern IC die are thinned to increase thermal conductivity and reduce package height. The target thickness for this version 0.2 design is 100 microns, but much thinner silicon wafers are used in current production, often loosely bonded to a thicker "handle" wafer for ease of processing. The server-sat will likely be built and tested with a thick handle wafer attached, but the handle wafer will be removed when the server-sat is attached to the deployment stack.

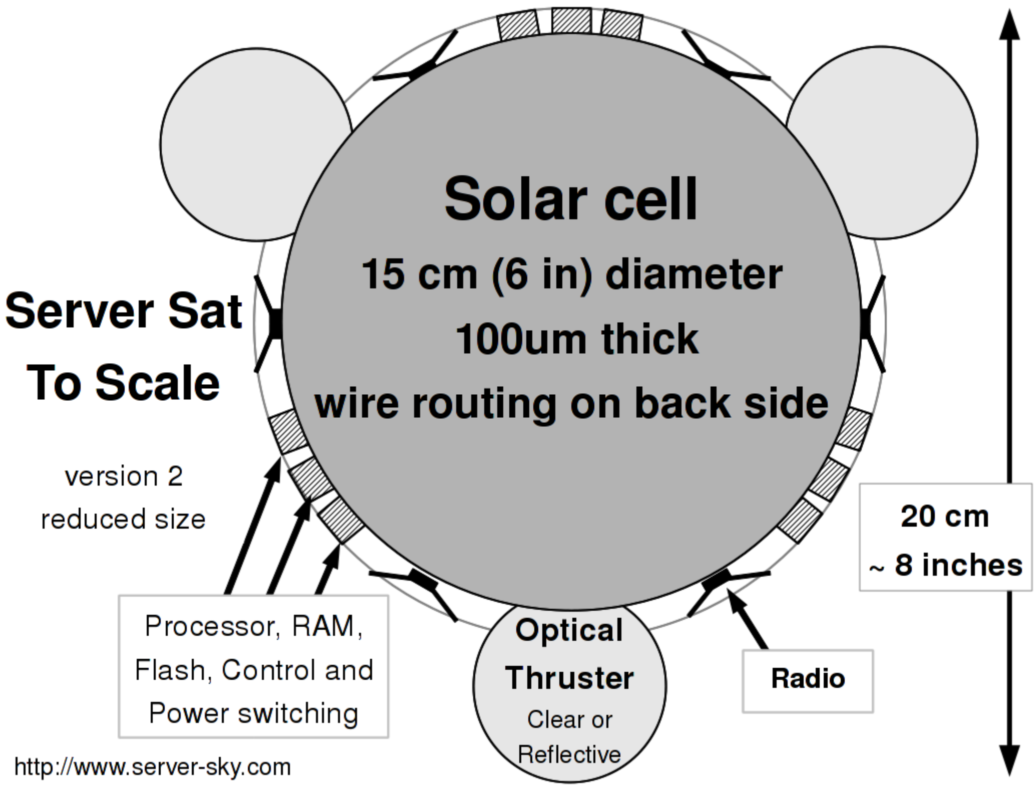

The above drawing is 50% of full scale. 10,000 server-sats can be stacked into a solid 1 meter column.

Reduced weight reduces launch cost and results in a more effective solar sail. The current thickness target is 100 microns, though production silicon is often thinned to as little as 20 microns for some applications. 100 micron thick silicon is very flexible, and can be rolled to diameters less than a centimeter without breaking. It is not unreasonable to assume that future server-sats can be as thin as 5 or 10 microns, weighing less than a gram.

Everything is coplanar - the chips are arranged around the edge, and power is fed outwards (in separated zones) from the solar cell. If a portion of the cell shorts out or is otherwise damaged, the remaining circuitry should still work.

If the solar cell at end-of-life is 13% efficient, then the solar cell will produce 6 watts.

Electronics: A server-sat used for database and web service may need as much as a terabit of flash memory (note - databases will be distributed over many server-sats). That will be about 4 by 4 centimeters of silicon area. Computational server-sats will need much less. While some high-performance processors and chip-sets use hundreds of watts, Giga-instruction-per-second level machines can get by with far less. For example, the PC Engines ALIX, based on the AMD X86 Geode processor, is a complete 4 watt system (including IO and power conversion losses) with 990 bogo-MIPS performance

[ALIX] . Optimized server-sats should be able to do far better.

Because the server-sat is extremely thin, some common electronic components cannot be used: electrolytic capacitors, cored inductors, etc. Bypass capacitors can be made thin, but it is better to keep peak impulse currents low. Some components such as crystal frequency standards may be replaced by surface acoustic wave (SAW) devices. However, other devices such as radios can be operated at low voltages and low impedances, and if some devices need higher voltages at trickles of current (such as LCD electrodes) they can be powered with capacitive charge pumps. At microwave frequencies, resonators can also be made with strip-lines and other components.

A server-sat will use a small array of radios (many more than the 6 shown) to communicate with neighbors in the array, with other arrays, and with the ground. Besides communication, server-sats will measure radio propagation time to neighbors to accurately compute spacing and orientation, with additional location information provided by other arrays, ground stations, and GPS. Multiple bands will be used, with frequencies that can penetrate atmosphere and clouds used for the down-links, and other atmosphere-opaque bands used for server-sat to server-sat communication. The server-sats will *not* have dishes, but will act together as a phased array antenna. Given the wide array spacing, there will be many spurious lobes, but it will still be possible to compute solutions allowing separate beams to many ground stations, far more than a traditional dish-and-transponder com-sat. While each server-sat may only dedicate one or two watts to the ground transmitters, the sum of thousands of transmitters will allow quite a lot of power for each beam. With the server-sat in a 4 hour orbit, it will be 7 times closer than a geosynchronous com-sat, so there will be a 50x advantage in beam power and ground spot area. Round trip ping time will be 70 milliseconds, less than U.S. transcontinental ping time through optical fiber.

Three optical thrusters are shown, which perform pitch and yaw control for the server-sat. Normally, the server-sat is pointed straight at the sun, and the thrusters are electrically stimulated into transparency. If one or two of the disks are unstimulated, they turn reflective, and produce about half a nanoNewton of thrust. This is enough to slowly rotate the disk. If all the thrusters are stimulated, this increases center-line thrust. Again, the thrust difference is not much, but it is enough to keep an array of server-sats on station relative to each other.

There is no direct way to control roll. However, the server-sat can be rolled by combinations of pitch and yaw [SSAT]
In orbit, the stresses are very small; the disk is relatively rigid by comparison. Maneuvers such as rotation and acceleration will take hours to days; micro-gee forces are involved. The main stresses on the disk will be from thermal contraction. When exposed to sunlight, the disk will be at 300K; as it passes into earth shadow, it will quickly cool to 150K or less, heated only by the infrared radiation from the night side of the earth. The system is mostly silicon, with some glass and metal. Interconnect metal layers should be designed with strain relief, and the whole metal and insulator stack should have about the same average thermal coefficient as the silicon.

Server-sats are designed to have uniform thickness and matched mechanical properties, so they can be stacked in cylinders for launch. Some kind of very thin separator may be needed, or the server-sats may stick together with vacuum welding or vander Waals forces. If a server-sat plus separators is 110 microns thick, then a stack of 10,000 server-sats will be 1.1 meters (43 inches) tall and weigh 70 kilograms (150 pounds).

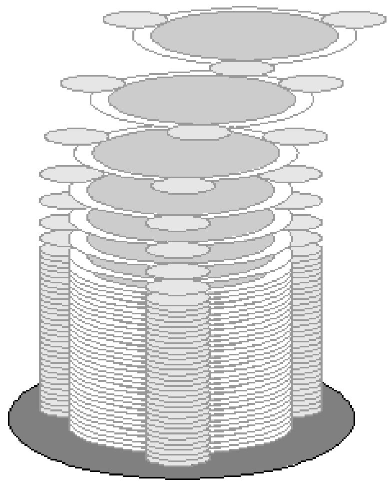

A typical modern geosynchronous communication satellite such as the HotBird 9 weighs 4880 kilograms (10,800 pounds) and has a 14kW array, launched by an Ariane 5 ECA, which can put 10,500 kg of satellite and apogee kick motor into a geosynchronous transfer orbit [HOTB] . The planned m288 server-sat orbit is lower; a larger payload is possible. 4200 kilograms is 600,000 server-sats and 3.6 megawatts of electricity. Thus, a server-sat array can produce more than 250 times the power (and communication capacity) of typical com-sats.

Optical Thrusters

Optical thrusters are two thin (30 micron) layers of commercial glass coated with of transparent Indium Tin Oxide (ITO) conductor on inner surfaces. The bottom layer glass is coated in separately controlled strips to permit partial functionality in spite of top to bottom shorts. In typical applications, a 1 micron gap is filled with electro-chromic shutter material and 1 micron diameter glass beads. A different spacing may be chosen if that improves performance or survivability.

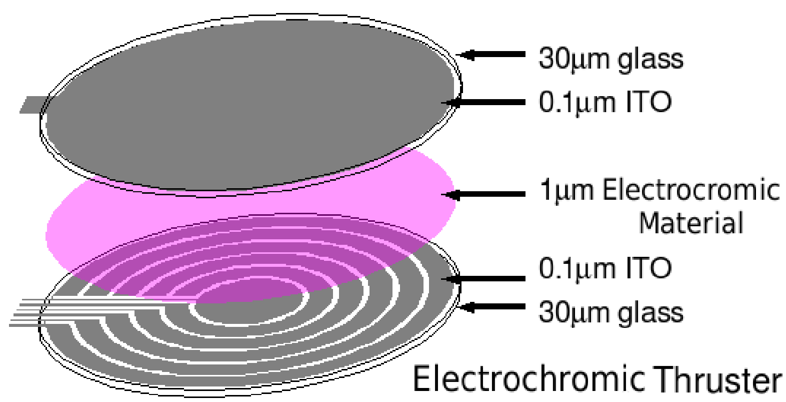

Liquid crystal light valves are well known, but they respond very slowly at low temperatures and do not have the best contrast ratio. A superior alternative may be electro-chromic materials, which are now being used as electronic window shades on the 787, and on whole buildings.

Maneuvering

A server-sat is light enough to be significantly accelerated by light pressure. At the earth's distance from the sun, the illumination is 1300 Watts per square meter, on average. The light pressure for absorbed light is the power divided by the speed of light, or about 4E-6 N/m^2 or 4 microPascal. If the light is reflected, the pressure doubles to 8 microPascal. This is a tiny pressure (sea level atmospheric pressure is 100 kiloPascals) but it is continuous. When pushing on something as thin and light as a server-sat, it can add significant velocity over hours, weeks, and years. The areal density of a 100 micron thick server-sat is 0.233kg/m^2 , and the albedo of a solar cell is around 0.15, so the acceleration is 1.15x4e-6/0.233 or approximately 20 micrometers/second^2, or 7 centimeters/minute^2, or 256 meters/hour^2. That allows for significant local maneuvering.

Large orbital changes are harder. Server-sats are in orbit, and if they are pointed directly at the sun, they are accelerated directly away from it. That adds to orbital velocity as their orbit takes them away from the sun, but subtracts from orbital velocity as they approach it. If they are tilted in relation to the sun, less area is exposed to light pressure, and the "albedo vector" of reflected light is tilted also, which can add a small sideways thrust.

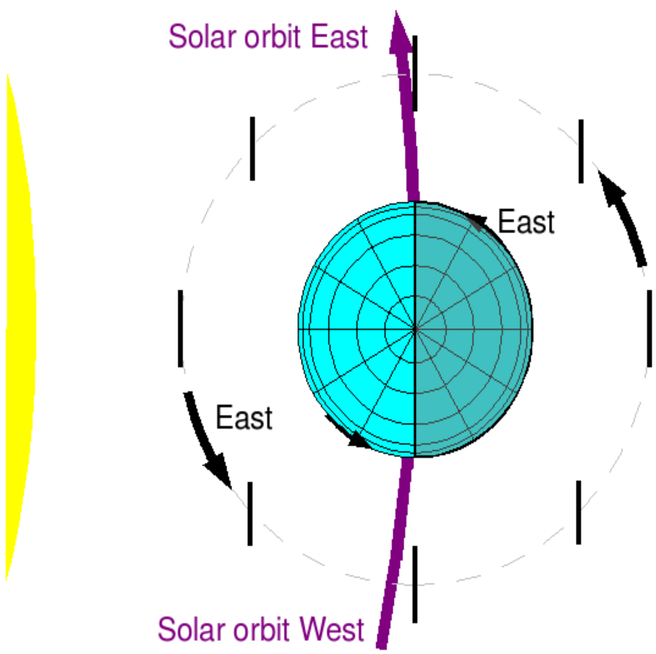

Defining directions

The earth rotates towards the east, counterclockwise when viewed from the north pole. It makes 366.24 turns relative to the fixed stars per year, and makes 365.24 turns relative to the sun. On the surface of the rotating earth at midnight, east is in the direction of the earth's orbit around the sun, and on the surface of the rotating earth at noon, the earth appears to be moving west. For the sake of argument, we will assume that the **earth is moving east in its orbit**. Thus, an object in orbit around the earth is moving towards the sun on the east side of the earth, and away from the sun on the west side of the earth. **I may have this backwards compared to some convention, so please add a note and a reference to the wiki if I goofed this up!**

Server-sats will get the most power if they face directly into the sun. However, they tilt to manuever. and that reduces the thrust. If they are tilted 45 degrees sideways, they get 30% less light and must reduce computing and radio functions, but they will still operate. With a 60 degree tilt, they get half power (and they cool down a lot!). Infrared light from the earth is mostly absorbed by the server-sat, and that creates some light pressure, too.

Light pressure from Optical thrusters

The version 1 design has three round optical-shutter light-pressure thrusters at 120 degree angles around the periphery. These are either reflecting or transparent. They are 5 cm in diameter (about 2 inches), and have areas of 2e-3 m^2. When reflecting, an ideal thruster produces perhaps 16 nano-Newtons, and when transparent it produces zero. Real materials will always show some reflection in transparent mode, and some transparency in reflective mode . Also, radiation from the earth (both albedo and infrared) reduces the effective thrust.

So, the thrust may vary between 12nN and 4nN (WAG). If one thruster on one side is fully reflective, while the other two are clear, the thrusters together produce a torque of 8nN times 10 cm or 800 pico-Newton-meters. If the entire server-sat has a mass of 0.007kg and an average diameter of 8 cm, the angular acceleration is 70 micro-radians per second squared. Accelerating for 36 seconds, then decelerating (applying opposite acceleration) for 36 seconds, will turn the array 10 degrees. Accelerating for 90 seconds, then decelerating for 90 seconds, turns the array approximately 60 degrees (not quite, as the thrusters are moving out of plane and become less effective when turned away from the sun).

Correcting for tidal forces

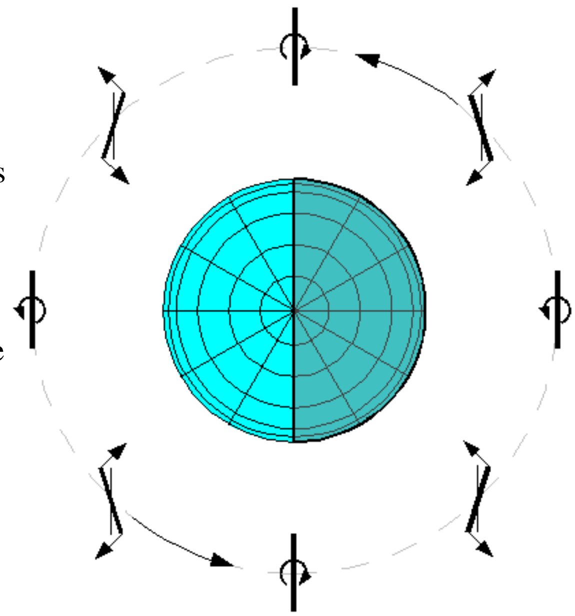

At the four "45 degree" points in the orbit, the server-sat is accelerated by tidal forces - the nearer end is pulled inwards by slightly more gravity and slightly less acceleration, and the farther end is pushed outward. These tidal forces are proportional to the vertical distance: $F=3\omega^2 M L \sin\theta$ where M is the effective mass at distance L from the center, $\omega=2\pi/Period$ is the angular frequency of the orbit, and $\delta=\omega t$ is the angle of the server-sat from the tangent of the orbit. The torque is proportional to the horizontal distance, or $T = FL\cos\delta = 3\omega^2 ML^2 \sin\delta\cos\delta = (3/2)\,\omega^2 ML^2 \sin 2\delta$.

The torque and the angular acceleration are maximized at a 45 degree angle. Both the torque and the moment of the server-sat are proportional to ML^2, so the angular acceleration is $\ddot{\theta} = (3/2)\,\omega^2 \sin 2\delta = (3/2)\omega^2 \sin 2\omega t$. This can be integrated twice to find the angular displacement from flat towards the sun: $\theta{=}{\approx}{-}(3/8)\sin 2\omega t$. The maximum angular displacement is given by $\theta_{max} \approx \pm 3/8 \text{radians} \approx \pm 21.5^\circ$ Although this is the "natural" oscillation if the server starts out flat, this is a metastable balance. Other perturbations such as tidal gradients from the sun and the moon will eventually displace the server into its lowest energy configuration, which is coplanar with the orbit. Hence, we will need at least some correction of the orientation.

Fortunately, the optical thrusters are much more powerful than the tidal forces, and can easily keep the server flat towards the sun. The maximum angular acceleration of the server-sat is $\ddot{\theta}_{max}=(3/2)\omega^2$ or 0.28 micro-radians per second squared for the m288 orbit, while the 5cm thrusters can provide angular accelerations of 70 micro-radians per second squared. This suggests a maximum mass-to-thruster ratio for server-sats: In the m288 orbit, the mass can grow to perhaps 50 grams, or somewhat more with larger thruster area percentage. Keep in mind that much larger thrusters will add more moment as well as more thrust. Since the thrusters will probably degrade over time, a reasonable safety factor is needed as well. In any case, centimeter-thick server-sats are probably out of the question in the m288 orbit, though they might be possible in m720 orbits, with half the angular frequency.

Thrust versus angle

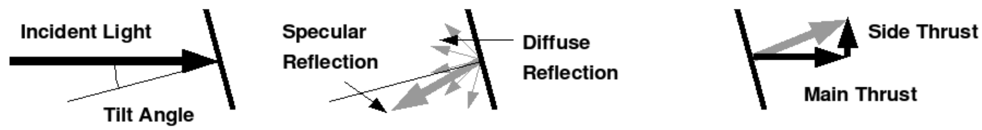

Most of the area of a server-sat is a big solar cell, which absorbs most of the light that hits it. Some portion of the light reflects from the solar cell, and the reflections can be roughly divided into **diffuse reflections** (in all directions from the front side) and **specular reflections** (opposite the incoming angle, like a good mirror). The diffuse reflections add an effective thrust of about 66% ($2/\pi$) of the diffuse reflected light pressure at the tilt angle of the server-sat, while the specular reflections add a thrust of 100% of the light at twice the tilt angle.

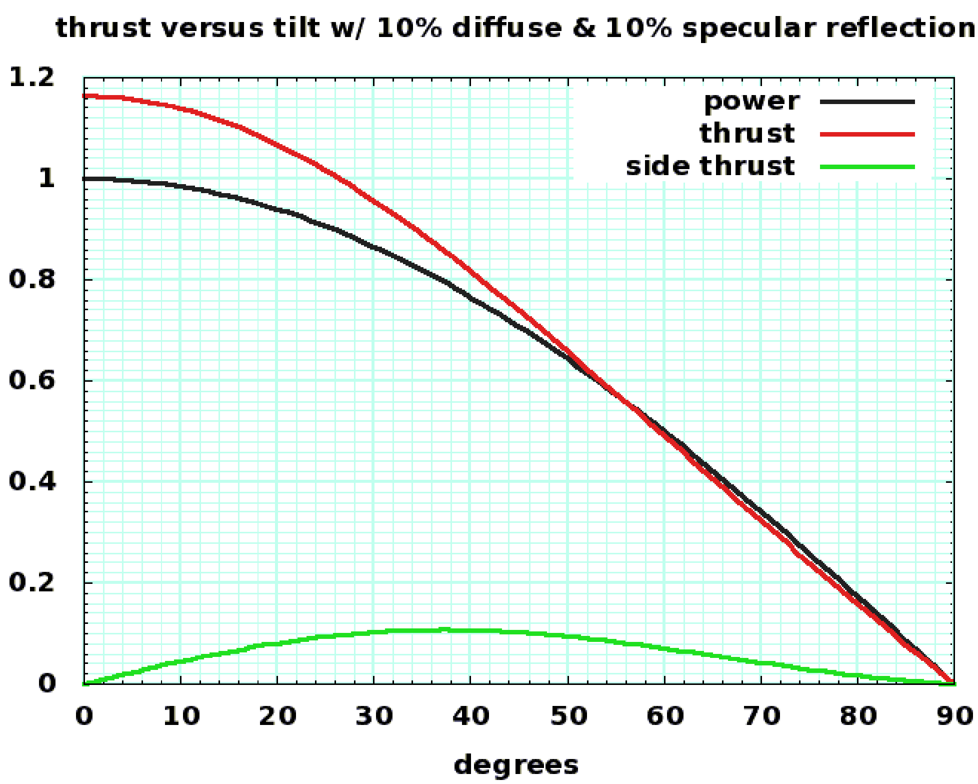

Here is a plot of the normalized main thrust (relative to of power divided by the speed of light, in the direction of the sun) and the side thrust (normal to the direction of the main thrust) as a function of tilt from the direction to the sun. The normalized power (relative to max power) is also shown.

Drag and Ballistic Coefficient

The first planned server-sat constellations are in the m288 orbit, at 6411 kilometer altitude. The atmosphere is very thin at that altitude, so drag will be negligible. However, server-sats (or fragments of them) may find themselves at lower altitudes, so the ballistic coefficient is needed to compute the decay of their orbits. The worst case ballistic coefficient is probably that of a flat plate moving face-on into the airstream, and the best case is edge-on. Lets assume the average resembles a sphere of the same radius (drag coefficient of 2), as it will probably be tumbling end over end. For a 9 cm disk weighing 7 grams, the ballistic coefficent is 0.15 kg/m^2, about 30% of the Echo communication balloon.

At the altitude of Teledesic and Globalstar, about 1400km, the mean atmospheric density is around 7.1E-15 kg/m^3. The decay rate at that altitude will be about 4000 meters/year; the velocity change needed to maintain orbit would be about 4 meters/second/year, or 0.12 microns/second/second. A server-sat can do that, with some maneuvering and perhaps some additional specular albedo added to the sun-side.

Deployment Orbits

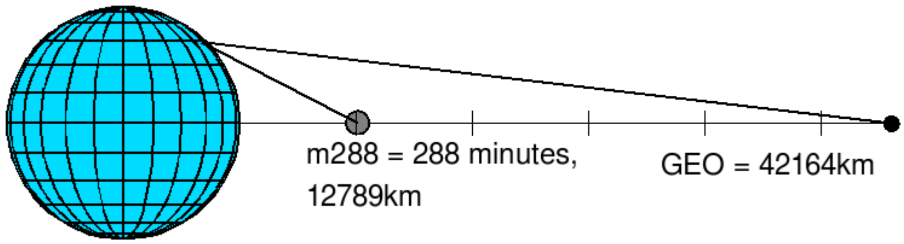

The first server-sat arrays will be deployed in a "4 hour" orbit, or more precisely a 23.9344696/6 = 3.989078 hour or 14360.7 second orbit (sidereal). This means it passes over the same spot on earth 5 times per day (= 6-1, the earth is turning underneath once per sidereal day), for a repeat time relative to the ground of 288 minutes (which we will call an m288 orbit). A 4 hour equatorial circular orbit has a radius of 12789 kilometers, and an altitude above the equator of 6411 kilometers. This puts it in a "thinner" part of the Van Allen belt, with an estimated unshielded radiation dose of 1Mrad/year [citation needed].

Relative to a position on the earth, a server-sat will be visible in the same position 5 times per solar day, at intervals of 288 minutes. The constellation of associated orbits will therefore be called the **m288** constellation. Of course, higher and lower orbit constellations are possible, though they will get more radiation. The m240 constellation will repeat 6 times per day, and the m360 constellation will repeat 4 times per day. For reasons that will become apparent, it is important that the orbital period is an integer fraction of a day.

The m288 central orbit can be seen at 58 degrees north and south latitude, at a distance of 10500 km. The round trip ping time is 70 milliseconds. The ground ping time through optical fiber across the United States is faster in theory, but ground networks are slowed by switches and indirect routes. Ping times from a fat-pipe server in Dallas Texas to mit.edu are 42 milliseconds, and to orst.edu are 49 milliseconds, so 70msec is not way out of line. However, much of the routing will travel "around the cloud", and without local caching in the "near" links, some pings may need as much as 200 milliseconds to hop from the far side of the orbit. Still, this is better than the 250+ millisecond ping time through a geosynchronous satellite.

Server-sats will actually be deployed in very slightly inclined and elliptical orbits, which map onto a nested set of tori ("torus-es") centered on the 4 hour, zero-inclination equatorial orbit. These are 4 hour orbits as well, and the largest tori are have major radii slightly offset inwards. A plane drawn perpendicularly to the center orbit (which intercepts the orbit in two places, and also intersects the center of the earth) will have "territories" marked on it for the regions around the various orbits, and the orbits passing through each "territory" can be treated as a property. This orbital property is further subdivided into angles around the orbit. Server-sats precisely positioned in each orbit will never intersect server-sats in other properties. Arrays in orbits in the same "property" will never intersect the orbits of arrays in different properties. This allows a large region of space to be filled with server-sats, potentially trillions of them. The density is limited by shading - at some point server-sats closer to the sun will reduce the daylight falling on the ones in the "back" of the orbit, and may begin to detectably reduce the sunlight falling on the earth.

Shading

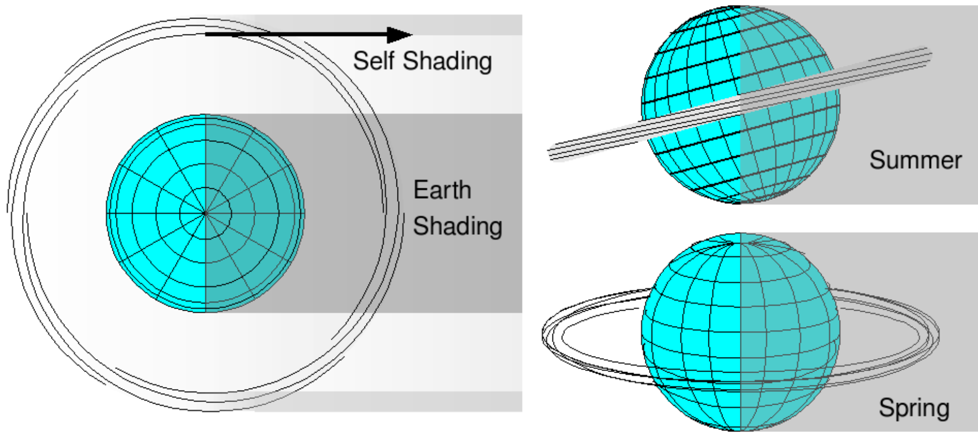

Server-sats will be spaced perhaps 50 meters apart in an array - an array with 32768 server-sats will be 1600 meters on a side. This puts them far enough apart that the shade area behind one 0.15 meter diameter server-sat will never completely block sunlight to the server-sat behind it. If the nested tori extend outwards to 500 km around the central orbit, that is a spatial volume of 6E10 cubic kilometers. Potentially, that is room for 500 trillion server-sats at an average 50 meter spacing. This "fuzzy toroidal cloud" of server-sats will block some sunlight, both to the server-sats in the back of the toroidal cloud, and to the surface of the earth. With 500 trillion 140mm diameter server-sats, the light blockage to the ground would be about 14% at noontime near the equator, with the blockage zone following winter as shown above. The blockage of back server-sats near the sides of the orbit would be more severe - at the equinoxes, it could be as much as 95% for portions of a fully populated array.

However, it is hard to imagine needing that many server-sats, even if they were operating mostly as space solar power satellites.   Beaming about 5 watts each to the ground, this far exceeds the projected world demand for electricity. With a "mere" 10 trillion server-sats, the blockages would be 2% to the ground and 38% to the back of the array. Also, for power production (and one-way information broadcast) ping time is no longer an issue, so m360 and higher orbits can be used.

Shading and power beaming will be partly compensatory in the amount of energy beamed to earth. Part of the sun's energy will be blocked, cooling the earth. Part of the energy to a wider band will be delivered at somewhat low efficiency as electricity (and later heat) to the surface. The global effects can probably be balanced, although to second order the equator will be cooled a bit, and the highly populated electricity-using temperate zones will be warmed a bit. That will probably affect weather patterns if it is done at an enormous scale. The amounts of power required to have a significant effect will probably be on the order of 200 Terawatts.

Nighttime Illumination

But shading is not the real problem. Because of the biological effects of accidental night-time reflection, server-sats should never be deployed in such densities this close to the ground.

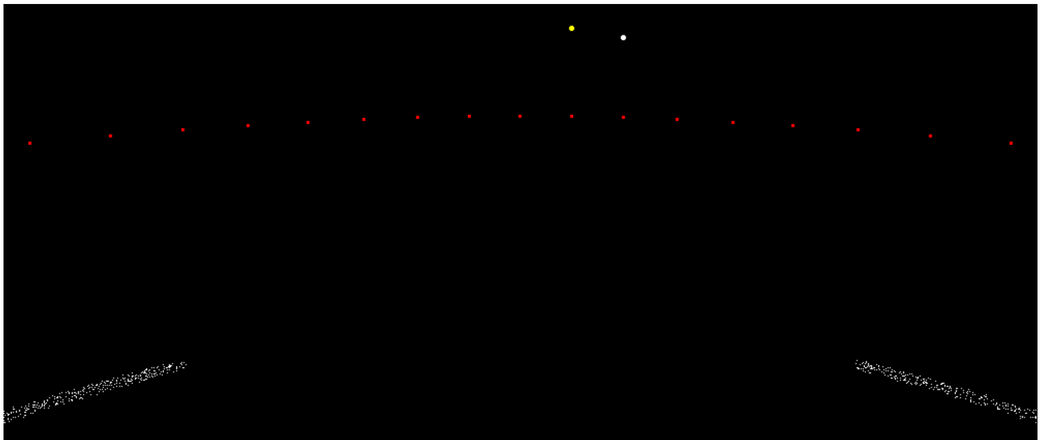

View of M288, geo, and moon

Server-sats will reflect some light. Oriented directly at the sun, most of the reflected light will go back to the sun, but some reflection will be diffuse and off angle. The sum of the diffuse reflected light from trillions of server-sats will appear as cloudy light in a band around the equator, interrupted by the Earth's shadow. At midnight on the equator, the illuminated bands will appear between the horizon (5.2 degrees wide) and 45 degrees from zenith (6.3 degrees wide). The two segments represent a total of 60 degrees of m288 orbital arc. The bands will be brightest just before the terminator, which will be slightly concave from the curvature of the earth's shadow.

How much light? For 10 trillion (1E13) server-sats, the total light will be about **10 times the total light from the full moon**. The moon has an albedo of 0.12, and so it reflects about 1.5E15 watts of light into a solid angle of 2π steradians. At 384000km away, we see about 1.6 milliwatts of light per meter^2 (approximately 0.3 lux). A server-sat will diffusely reflect perhaps 5% of the 130 watts of sunlight hitting it - 6.5 watts per server-sat. The visible 1/6 of 10 trillion server-sats would reflect perhaps 10 trillion watts of light, also into a solid angle of 2π steradians. At an average of 10,000 kilometers distance, the surface at the equator gets about 16 mW/m^2, or approximately 3 lux. Since recommended office illumination is about 300 lux, that is way too much nighttime illumination!

Without much better control of reflection, the total number of server-sats in at m288 will be limited by light pollution and its effects on nocturnal animals (for example, corals are sensitive to the monthly lunar illumination cycle, probably to a fraction of full moon light. **This may limit the near-earth constellation to less than 1 trillion server-sats**. However, by actively orienting server-sats so the reflection is directed away from the earth, and by reducing off-axis reflection (which is also good for thrust control), the amount of reflected light reaching the earth can be greatly reduced. This will need more study.

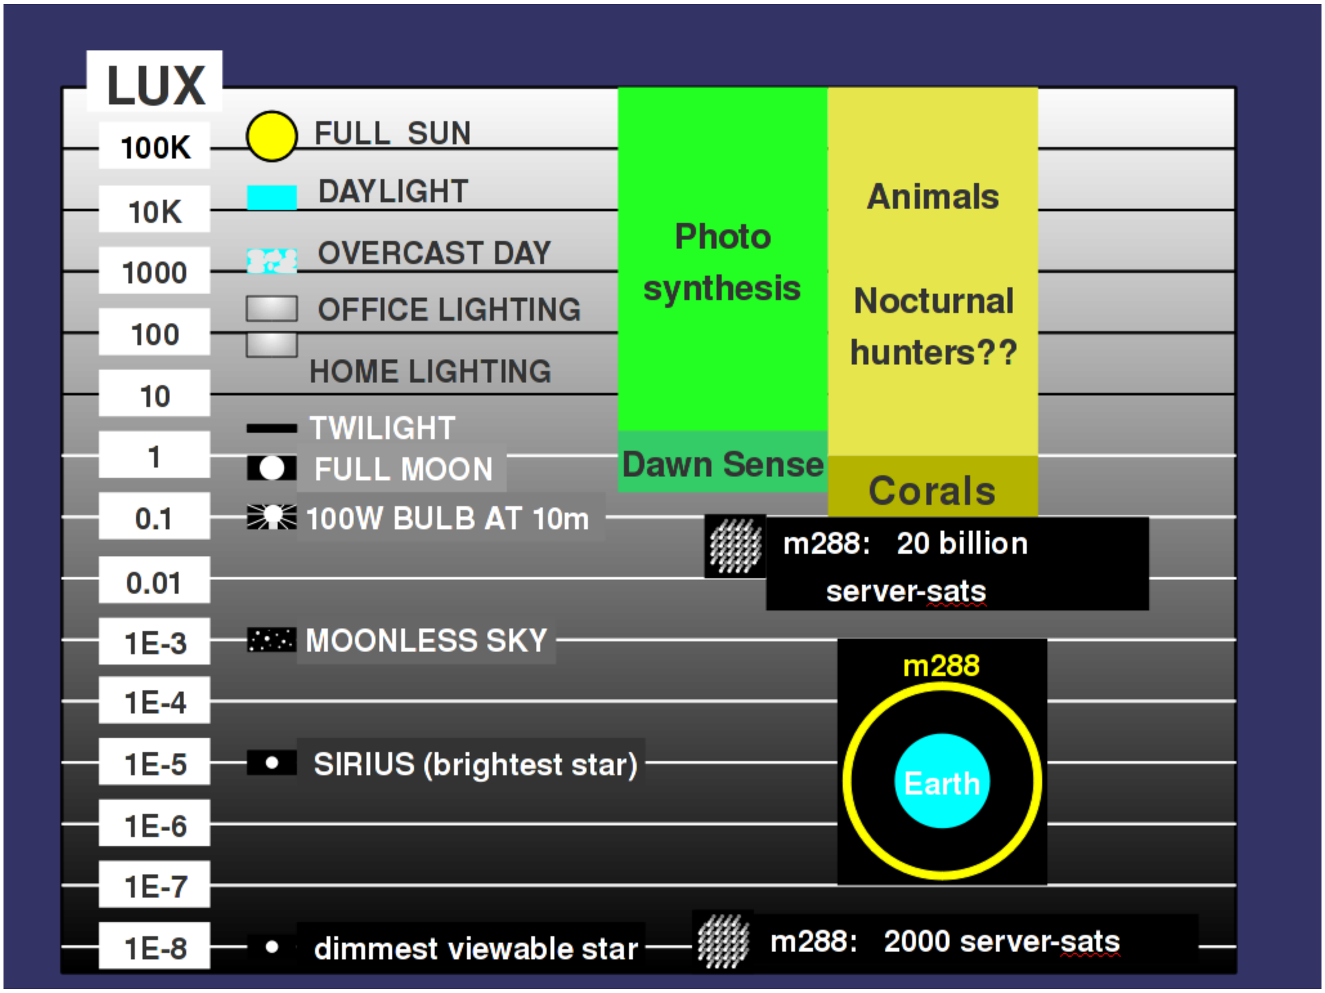

If 10 trillion server-sats are deployed as far out as the earth-moon L4 and L5 Lagrange points, the situation is much better. The constellations will be visible as often as the moon, though 60 degrees offset to east or west. On average, each constellation of 5E12 server-sats will be reflecting 3E13 watts, and at the same distance as the moon, each will appear to be about 2% of the brightness of the full moon, much less with reflectivity improvements and attitude control. If server-sats are deployed for large-scale power generation, where ping time is not important, farther out is better. However, two constellations at L5 and L4 will not be able to send power to the earth surface on the opposite side of the moon. A third constellation at the metastable L3 Lagrange point will also be needed to provide full ground coverage, and that will require constant station-keeping to stay in place.

Elevation above the horizon, interference with geosynchronous satellites

Assume that the first torus to be filled with arrays is at a minor radius of 10 km from the central orbit. Refraction will lift the apparent elevation of the central orbit a few degrees above the horizon, so a south-facing antenna at 60 degrees north or south may be able to see the edges of the torus. However, ground sites further north or south of the central band may need to relay through other satellites, or through landlines on the ground. Between latitudes 10 degrees north and south, the torus will have similiar ground antenna elevation angles as geosynchronous satellites. Assuming that the same satellite frequency bands are used for server-sky as for existing geosynchronous services, latitudes between 10 degrees and 60 degrees north and 10 degrees and 60 degrees south should be able to make direct use of server-sky.

Intentional gaps in the constellation

The orbits will not be completely filled - there will be large gaps in them to permit transit of launch vehicles. Establishing these "windows" will involve much negotiation, beyond the scope of this document.

If it ever proves practical to build space elevators, there will need to be other gaps in time and "property" so that the elevators and the server-sats never interact. Fragments of a failing space elevator will collide with an undetermined number of server-sats (as well as colliding with and destroying all the other space elevators), so space elevators and the permanent use of these orbits (or possibly any near-earth orbit) may prove incompatible.

In the short term, there will be millions, not trillions, of server-sats. They should still be assigned positions and orbits compatible with a much more crowded sky. It is good to know, going into this, that the region available for well behave server-sky orbits offers much room for growth.

Defining Property Within the m288 toroid

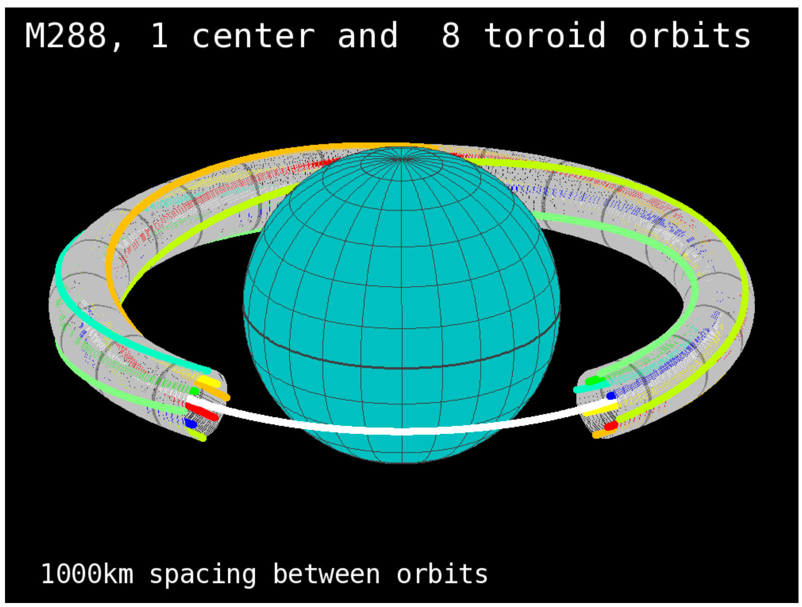

All server sky orbits in the m288 toroid will have the same identical semi-major axis, and hence the same identical 288 minute synoptic period, and the same 14393 second sidereal period. If the orbits are mapped correctly, every item in them will maintain the same approximate spacing to neighbors in three dimensions, even as the whole constellation makes one orbit around the earth and one "short axis rotation" around the central orbit.

This allows very large numbers of server-sats to be deployed in the 60 billion cubic kilometers of server sky. Properties can be assigned much like IPV6 address space is assigned by ICANN - indeed, we may map IPV6 addresses onto particular orbital volumes, and know where in the sky a particular server-sat is from its IPV6 address, or vice versa. If we map down to one meter cubes, we will need about 65 bits of address space.

The space does not map in a rigidly cartesian way; the "toroids" we are mapping on are slightly distorted from perfectly circular minor cross sections, as shown:

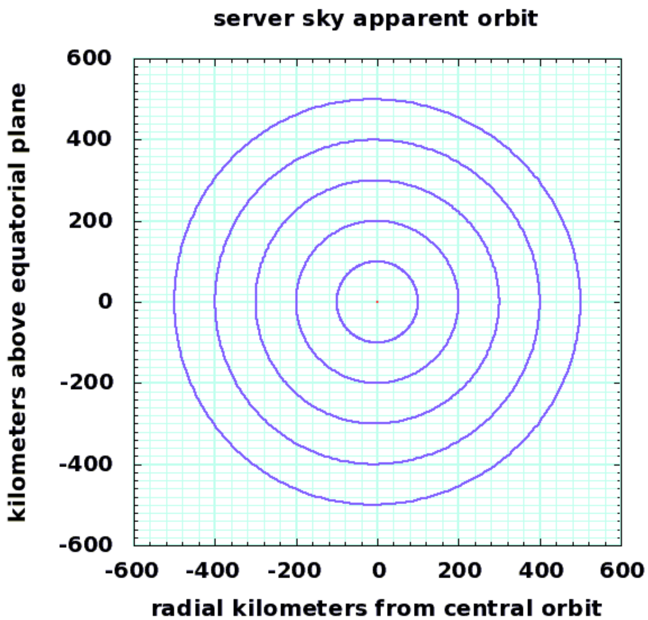

500km cross section of nested torii around the m288, 12789km central orbit region. The outer orbits shown have an eccentricity of 0.039 and an inclination of 2.24 degrees. Although these torus cross sections look circular, they are actually distorted by as much as 3 kilometers from perfect circles. The outer torus has a volume of about 60 billion cubic kilometers. The inner torus, with a radius of 100km, has a volume of 2.5 billion cubic kilometers.

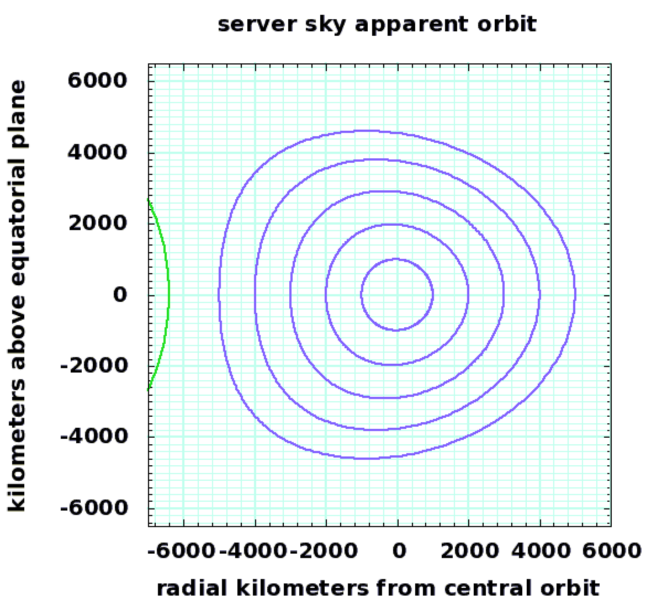

5000km cross section of nested torii around the m288, 12789km central orbit region. The outer orbits shown have an eccentricity of 0.39 and an inclination of 23 degrees. The edge of the earth is to the left. These unusually broad toruses demonstrate the distortion that occurs in outer orbits. The outer torus has a volume of about 6 trillion cubic kilometers (the earth is 1.09 trillion cubic kilometers), a perigee of 1411 kilometers altitude, and intercepts the orbits of many existing satellites and large debris objects, so such large torii should not actually be used.

In addition, objects in the orbits will be advanced or delayed because of increased velocity at perigee and decreased velocity at apogee. This tends to spread objects near perigee (in the line of the orbit) by 1+eccentricity, and near apogee by 1-eccentricity. Thus, the volume elements will lengthen and shorten, and distort trapezoidally (proportional to the eccentricity, or the toroidal minor radius divided by the major radius) as they travel around the orbit. The arrays within those volumes will distort proportionally.

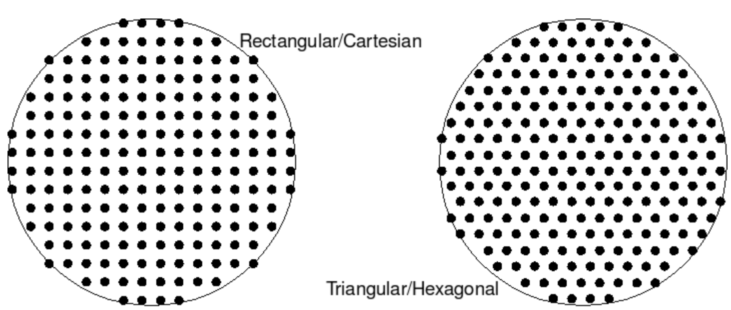

This is based on a rectangular cartesian grid mapped into a circle at one angle from the equator to the fixed stars - the prime meridian seems like a good candidate. Other grids, such as triangular/hexagonal, can be mapped within the same circle (a cross-section of the largest torus).

Assigning orbits within this map can be done many ways - by an organization like ICANN, perhaps working outwards from the center in the order of assignment.

An intriguing alternative is to map the surface of the globe onto a circle, and then assign orbits to the countries that fall on that mapping. That leaves open spaces in between, in the areas of the globe that are ocean (and on the map) for navigation in space. Here is one such mapping, an [Ernst Hammer equal area projection.]() It has been modified from an ellipse to a circle. Other possibilities include the [Mollweide projection.]() A more abstract projection that exaggerates continental areas, regularizes passages through the orbit torus, etc. is preferable.

Arrays in toroidal orbits will distort, because different parts of the array are moving at slightly different speeds. Server-sats moving towards apogee are skewed towards it, by approximately twice the array dimensions. This is called apogee-skew and is illustrated below (see the website for an animation):

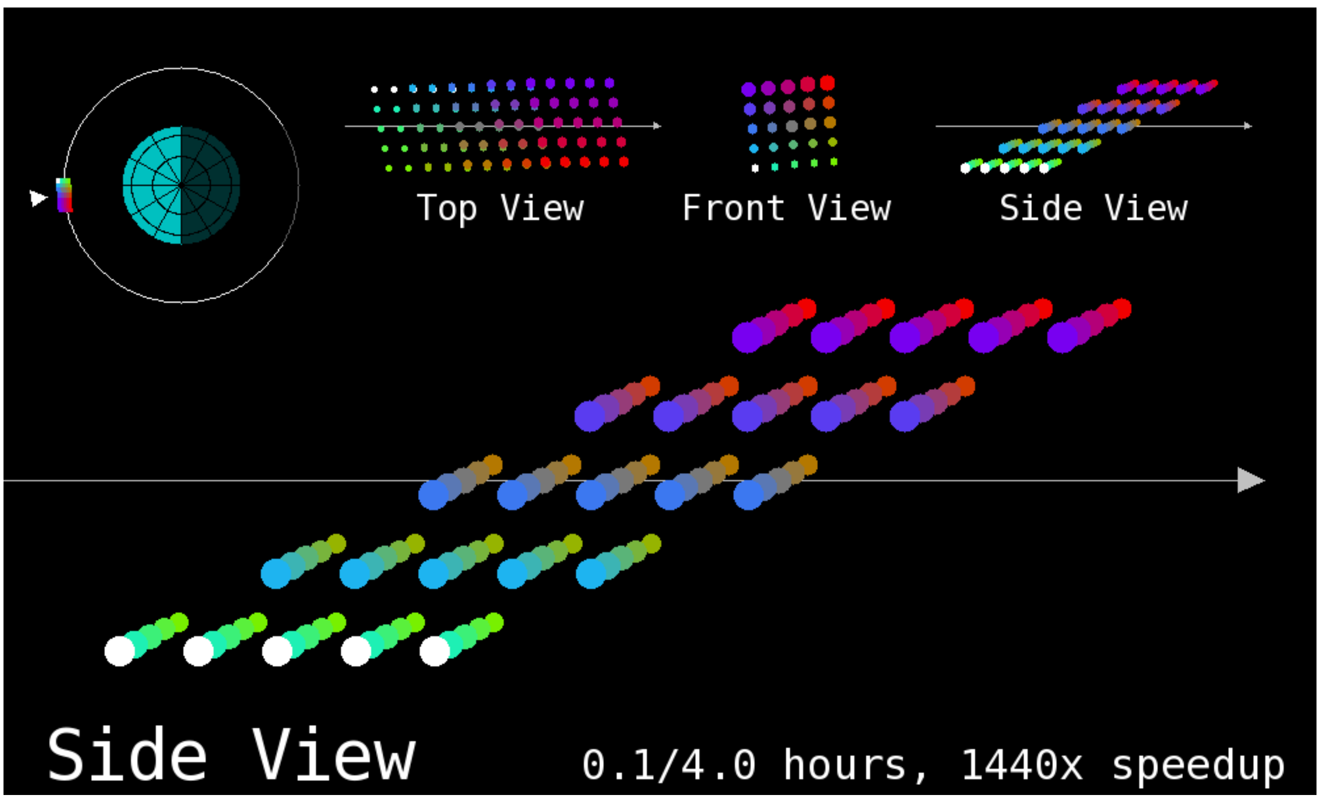

Launching Server-sats

Server-sats are deployed in an approximately equatorial orbit, so that the ground antennas can have a fixed elevation above the horizon. An inclined m288 orbit would change elevation up and down 5 times per day.

Earth orbit

									longitude	Kourou 5	KSC 30	Baikonur 46
orbit name	**surface period**	**lat45 ping**	**radius**	**radius**	**altitude**	**gravity**	**sidereal velocity**	**sidereal period**	**launch velocity**	**insertion velocity**	**insertion velocity**	**insertion velocity**
	minutes	ms	earths	km	km	m/s^2	m/s	sec	m/s	m/s	m/s	m/s
LEO	97	33	1.05	6678	300	8.94	7726	5431	7531	676	3847	6003
m288	288	63	2.01	12789	6411	2.44	5583	14393	8667	1118	2726	4072
m360	360	73	2.26	14441	8063	1.91	5254	17271	8846	1211	2592	4511
m480	480	87	2.63	16756	10378	1.42	4877	21585	9050	1308	2451	3516
m720	720	110	3.18	20295	13917	0.97	4432	28774	9287	1404	2297	3188
GEO	fixed	253	6.61	42164	35786	0.22	3075	86164	9954	1511	1860	2281

The mXXX series of orbits were chosen as even sub-multiples of a day. This facilitates coordination with launch system logistics (the same launch windows open up at the same times of day, simplifying staffing), and also simplifies interaction with space elevators, should those ever become practical.

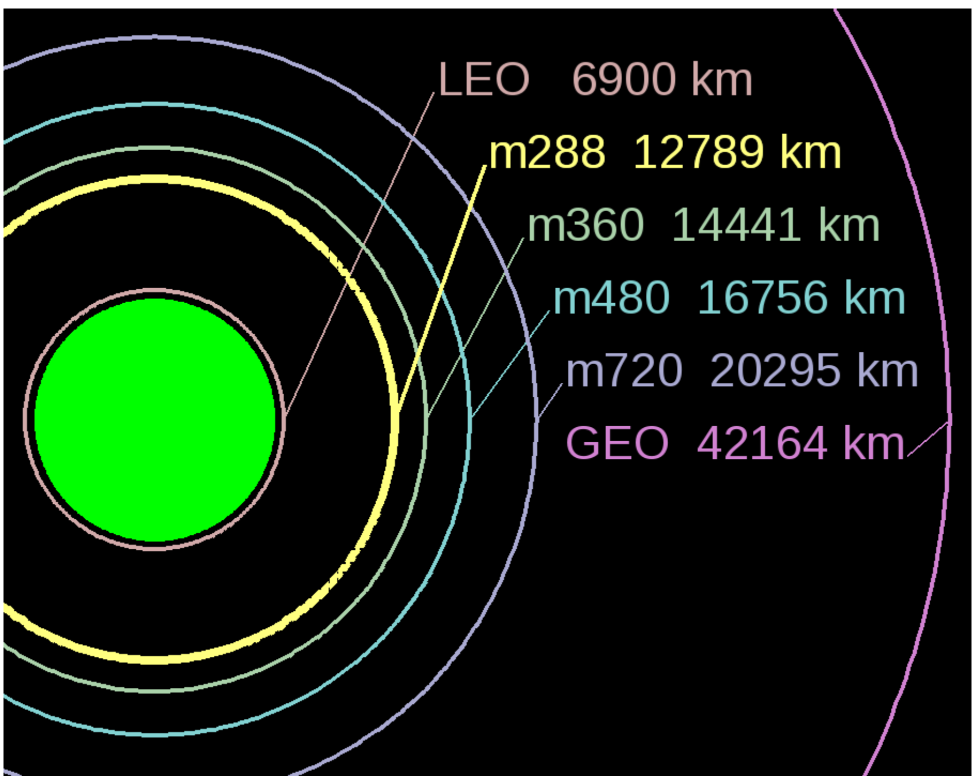

m288 was chosen as a compromise between visibility, launch cost, and ping time. Higher orbits will have more ground coverage, and spend a smaller fraction of their orbit in shadow, but will be more expensive to reach, and receive higher radiation doses in the Van Allen belt. Low-cost alternative launch systems may make higher orbits much easier to reach with much larger arrays, and change the optimum altitude.

While the current plan is to use toroidal orbits near the equatorial plane, this does *not* rule out polar orbits. If the orbital period is precisely controlled, with overlapping array gaps, then arrays of server-sats can be orbiting in both directions, providing coverage to high latitudes. This may also minimize night sky pollution. However, big problems may occur if the system is abandoned or goes out of control.

Radio Arrays

Server-sat radios will have multiple low-power outputs and communicate to many printed-circuit antennas and resonant impedance matching structures. They will talk on multiple bands, for downlink, uplink, femtosecond-precision array timing, micron-precision server-sat location and orientation within the array, and orientation to other arrays and to GPS and ground systems. Much of this accuracy will come from continuous monitoring and averaging, differential and quadrature analog signal processing, and the ultra-low vibration and perturbation environment of a server-sat in a completely predictable nano-gee space environment.

Timing precision and synchronization

In the late 1990s, Keith Lofstrom worked with Teradyne and Analog Devices to produce a timing interpolator system for a tester that produced timing edges with 1ps resolution within a 2ns time period. Measurements showed that the timing accuracy of each edge from the interpolator chip was less than the 8 femto-seconds resolution of the measurement equipment. This was for single edges - an system that averages trillions of periodic edges into planar resonators should be able to estimate system phasing to fractions of a femto-second;   at the speed of light, this is a small fraction of a micron. To achieve this, the measurement system must be well isolated from all noise sources, and all signals must be differential, with balanced power on all edges, balanced grounds and power signals to each circuit, etc.

Vibration can generate errors, but the vibrations will be very small in a server-sat (perhaps from nano-meteor impact) and measurable. Thermal variations will occur as various electronic sections turn on and off, but these thermal variations can be characterized and extracted. Perhaps the largest source of unpredictable perturbations will be radio energy from other server-sats coupling with unshielded signal runs on the server-sat; this may require extra ground planes.

Down-link communications

As a new service, server sky will likely be allocated EHF frequencies {EHF] for the down-link - in the 30GHz to 300GHz range. For now, let's assume a frequency of 38GHz and a wavelength of 8 millimeters. This wavelength is much smaller than a server-sat, so directional beams can be made with server-sat scale antenna arrays. Each server-sat can direct radio energy into an angle of perhaps $\sin^{-1}(0.1)$ or 6 degrees, for a ground spot of perhaps 600km.

However, the directionality of sever-sats comes from their ability to act as a phased array. Constructive and destructive interference between phase locked arrays of server-sats permits ground spots of a few tens of meters - better than cellular service and WIMAX. The wider the array, the smaller the ground spot, so for down-link at least, adding server-sats will improve spatial multiplexing bandwidth, with no practical limits on download bandwidth to billions of customers on earth.

A phased array works by adjusting the time delay of each server-sat transmitter so that the signals from each transmitter, located at a different distance from the receiver, all arrive at the receiver at same time. If each transmitter is emitting a pure sine wave, this can be accomplished by shifting the phase of the outgoing signal.

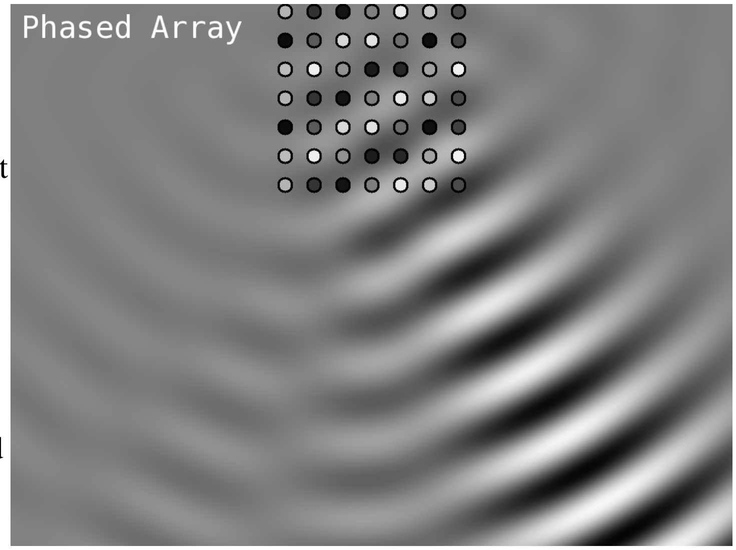

The easiest way to do this is to compute the path length from each transmitter to the ground receiver in wavelengths, take the fractional part, and conjugate it (that is, the negative fractional part becomes the phase of that transmitter). For a 10,000km path, that can be accurately represented as a 32.10 bit fixed point number or a 64 bit IEEE754 floating point number.

If the system is linear, then the transmitter can emit the sum of many different phased signals pointing at many ground spots. An easy way to do this is to make the transmitter emit modulated I and Q signals (90 degrees apart), where each I and Q signal is the computed algebraic sum of many base-band or intermediate frequency signals representing different spatial channels. Modern VLSI integrated circuits can combine many data channels and compute the phased sums of them at high speed, while recomputing transmit angles to accommodate the movement of the orbiting array relative to the ground (angles will change 21 nano-degrees per microsecond). In this way, one phased array can communicate with many different ground spots.

Three Dimensional Phased Arrays

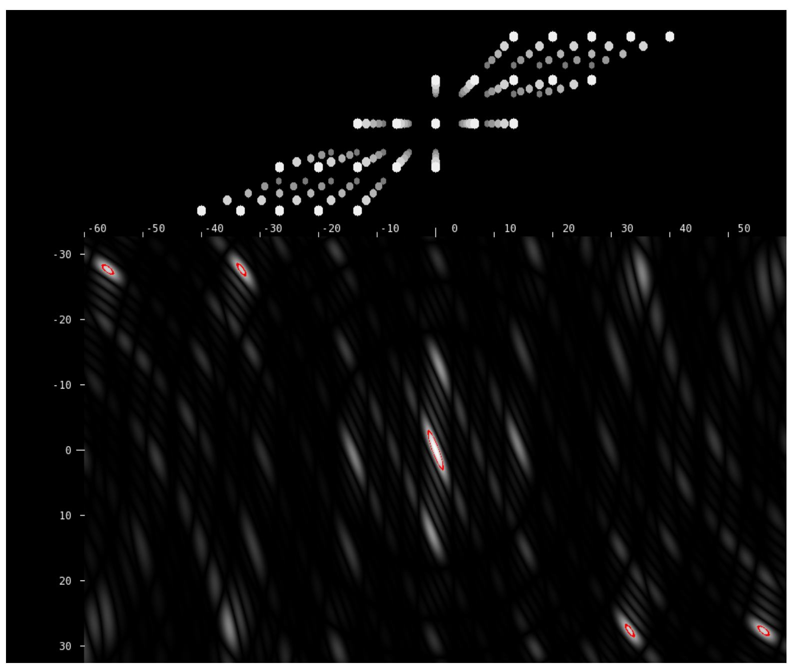

Server Sky server-sat arrays are widely spaced relative to the radio wavelength. However, the spacing in the three dimensional array can be continuously adjusted in 3 dimensions and 3 rotations, subject only to avoidance of radio and solar shading. The arrays move in relation to the target, and rotate around the central orbit, so the phasing changes continuously. However, it changes slowly compared to round-trip ping time.

The array spacings are large compared to a radio wavelength, so they may scatter side-lobe energy into grating lobes. Here is a ground power plot showing the grating lobes. The website shows an animation of the array rotating as it orbits, and the movement and shift of the bright spots, which represent a concentration of radio power

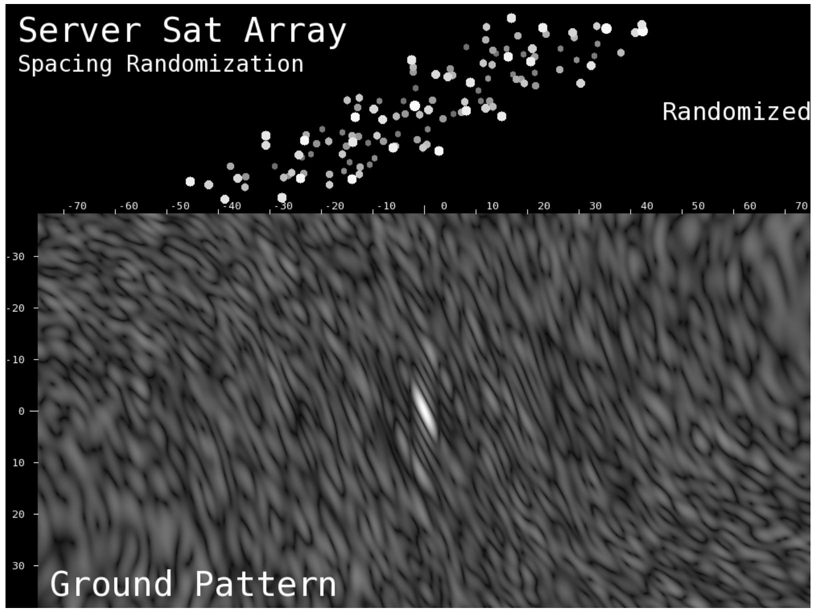

The brightest grating lobes can be smeared out and flattened by randomizing the spacing of the server-sats in the array. This has the same effect as a sparsely populated array of much more closely spaced server-sats. A much larger array will suppress the “clutter” noise by the square root of the number of server-sats.

While the precision spacing of the server-sats permits accurate random spacing, that is probably not the best way to reduce grating lobes. It does show that we can do better than a regularly spaced lattice, and further study will find many sets of functions that will better reduce grating lobes.

Keep in mind that each server-sat has many radio antennas located at fractional wavelength spacings. This means that each server-sat can do some beam-forming on its own. The wider angle grating lobes will receive less energy because of this. And these patterns emerge from continuous wave single frequency broadcasting - it is likely that side-lobe modulation, FM modulation, and code division multiplexing can all be used to minimize the interfering effects of grating lobes. If the server-sats have sufficiently agile and controllable radios, their localization performance may be greatly improved while they are in orbit, through spacing and radio software changes alone.

Communication within an array

Server sat arrays will communicate from satellite to satellite with individual high bandwidth 60GHz beams. All the satellites within an array will need precision timing information, and will also be getting copies of receive and transmit packets for computing phased array beams.

Communication between arrays

This will be difficult - not technically, but primarily due to licensing. The server-sky orbits will be approximately in-plane with equatorial communication satellites, which also use in-plane methods to communicate. This means that beams between server-sky arrays will continue beyond the server-sky orbit and up to GEO, causing interference with receivers there.

This can be partially ameliorated by using very narrow, oblique, and fractional-orbit beams between inclined server-sat arrays, such that the beams are pointed well above or below the orbital plane, and relatively weak when they reach GEO altitudes anyway. Space-to-space communications often use the

60GHz band, where the atmospheric absorption of oxygen is very high, which will isolate the space links from potential jamming on the ground. Here is an example of some oblique beams between arrays:

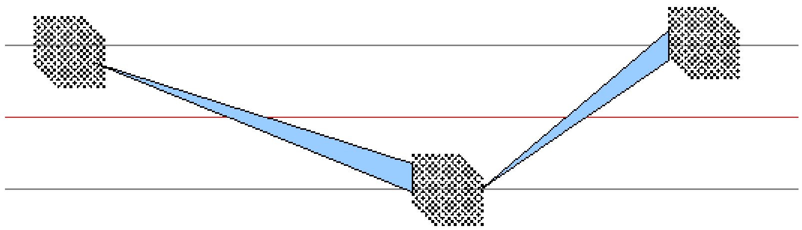

If the beams must miss the GEO orbit by, say, 5 degrees, and the communication partners are 300km above and below the equatorial plane, then they can be spaced 7000 km apart. This would require 7 hops to travel halfway around the ring (41000km). Actually, they should be spaced closer anyway, to reduce power. The speed-of-light propagation time around the ring is 136ms. If packets are 1400 bits and running 10Gbps, then a packet time is 140nsec. Assume that switching, re-route, queueing and beam forming time add a latency of 5 microseconds to each relay. 1000 hops would add only 5 milliseconds or 4% to the path latency, permitting an array spacing of 41 km, with the arrays only 2km above and below the plane.

As the arrays get denser and larger, the beam size gets smaller, but the amount of inter-array traffic will probably increase faster than the number of communication paths. This needs to be characterized.

Communication via other satellite services

Server-sky orbits near the equatorial plane. The constellation is visible only to ground sites with latitudes below 55 degrees or so, precluding much of northern Europe. While it is possible to add high inclination orbits in synchronization with the main constellation to reach far north, in the short term it is easier to use existing services like Iridium, Globalstar, TDRSS, and the many satellites in GEO to do this. In the near-term, server-sky is primarily doing computation and radar sensing, and the constellations around the equator will be incomplete anyway. It is better to rely on existing infrastructure when geometries and transponder configuration on the existing satellites permits this. This needs more study.

Radar - locating space debris

An important function of a server sky array will be locating space debris and other satellites. Server sat antennas are too noisy and non-directional to make good radar receivers. However, they make dandy transmitters; working in conjunction with existing radar satellites and ground stations, they can through very tight beams with high power density through space, and off-angle reflections off small bits of space debris can be detected by many radar receivers optimized for the purpose. This should permit much more accurate location and characterization of much smaller bits of space debris.

Orientation to other arrays, GPS, and ground stations

Server sky is blind. It does not have star trackers or ring laser gyros or any other typical orientation device. It may have some MEMs gyros and accelerometers, but those are fragile and expensive to develop and the thinning needed may cost too much for Commercial Off The Shelf (COTS) devices.

Server sky has two senses, though. Sunlight to the solar cell can be measured and used to determine the sun angle. Gratings can be added to sense sun direction. So a server-sat will have some limited optical orientation capability.

Server-sky's main sense is radio. It will be in constant communication with neighboring server-sats, and can do precision orientation and location computations from that. It can also sense signals from ground stations and GPS. Modulated radio and sub-wavelength fringes are used in commercial surveying equipment to measure distances with high precision; server-sky will do the same thing with the 60GHz intra-array communication links. Assuming it can measure phase within 1 degree at 60GHz, a server-sat should be able to locate its many antennas with a 20 micron accuracy. Better measurement capabilities can be added in the future by improving the software.

Server Sat Materials and Manufacturing

Devices and materials

A lot can be done on a very thin planar surface. Other things cannot be done easily. Here are some common electronic devices:

Thin Planar	**Non-planar**
Printed circuit laminates (ultrathin)	Connectors
Resistors	cooling fins
Planar Capacitors	Wound foil capacitors
pixel sensor arrays	lenses
electrochromic light valves	LEDs (?)
Strip-lines	Coax
Microwave-frequency inductors	Low frequency inductors and transformers
surface acoustic wave (SAW) resonators	crystals
beam lead interconnect	wirebonds and solderbumps

Silicon is the construction material of choice - the solar cell is made of silicon, and the processors and memory are also. Here are some relevant properties of silicon, SiO2 glass, gallium arsenide, copper, aluminum, silver, gold, tantalum, indium tin oxide, kovar, invar, low expansion borosilicate glass, and pyrolytic carbon, which will make up 99.9% of the weight of a server-sat:

property	Si	SiO2	Si3N4	GaAs	Cu	Al	Ag	Au	Ta	ITO	Kovar	Invar	BSiO2	PyroC

Density g/cm^3	2.33	2.65	3.20	5.32	8.96	2.7	10.5	19.3	16.7	6.43	7.85	8.05	2.23	2.25
Coeff. of Thermal Expansion 10-6/K	2.6	0.5	3.2	5.7	16.5	23.1	18.9	14.2	6.3	10?	5.3	1.3	3.25	4.3
Heat Capacity J/g-K	0.71	0.74	0.71	0.33	0.38	0.90	0.24	0.13	0.14	-	0.44	0.51	0.75	0.72
Heat Capacity MJ/m^3-K	1.65	1.96	2.27	1.76	3.40	2.43	2.47	2.49	2.34	-	3.45	4.11	1.67	1.62
Thermal Conductivity W/m-K	149	1	30	55	401	237	430	320	58	-	17.3	10.1	1.1	1950
Specific thermal conductivity	64	0.4	9.4	10	45	88	41	16	3.5	-	2.5	12	0.7	870
Thermal Diffusivity mm^2/s	90	0.5	13	31	118	98	174	129	25	-	5	2.5	0.3	1200
Youngs Modulus GPa	150	73	260	86	110	70	83	78	186	116	140	148	63	4.8
speed of sound km/s	8.0	5.2	9.0	4.0	3.5	5.1	2.8	2.0	3.3	4.2	4.2	4.3	5.3	1.5
elastic impedance Kg/mm^2-s	18.7	13.9	28.8	21.4	31.4	13.7	29.5	38.8	55.7	27.3	33.2	34.5	11.9	3.3
Tensile Strength MPa	7000	50	70	57	210	40	170	100	200	120	270	680	35	-
Atomic Weight (avg/atom)	28	20	20	72	64	27	108	197	181	57	57	56	19	6
Resistivity nano-ohm-m	-	-	-	-	17	27	16	22	131	2200	490	820	-	14000
Dielectric Constant	11.8	3.9	7.5	12.9	-	-	-	-	-	-	-	-	4.6	-

Data mostly from wikipedia and various places online. See also Matweb [MATW] the material properties website. The B.Y.U. CTE table [BYUC] is useful. The tensile strength numbers above untrustworthy, and many parameters are anisotropic. Use this table only for rough estimates.

The vast bulk of the material , and the largest pieces of of the server-sat, will be silicon. Since the server-sat undergoes wide temperature changes when it passes in and out of shadow, or undegoes thermal annealing, it will be more survivable if the non-silicon portions are made of composite materials that match silicon's 2.6E-6/K coefficient of thermal expansion (CTE).

Server sats will also need transparent materials and conductors that closely match silicon. The metals have very high CTEs, while SiO_2 has a very low CTE, so slotted metal wires with SiO2 in the gaps is one way to make a "material" that is both conductive and has the same CTE as silicon.

Stack compression during launch

Booster systems vibrate during launch. If there are regions of higher and lower compressibility and mass density, there will be standing waves and resonances in a stack of server-sats. Ideally, the design would match the mechanical properties of all the materials used, but that is unlikely given other constraints. So the stacks may need to be resonance isolated from the boosters, reducing the payload fraction. A somewhat simpler constraint is to match the mechanical compression of the stack caused by acceleration forces. The materials (plus spacers if necessary) should have the same ratio of compressibility (modulus) to mass density, that is, the same speed of sound. That will minimize shear forces on the connections from the solar cells to the electronics and thruster ring. Mismatches will restrict the number of server-sats that can be stacked between spacers. Of course, all these problems must be designed out with mechanical CAD, and tested with centrifuges and shock tables on the ground.

Composite Silicon and Pyrolytic Carbon matched to Borosilicate Glass

One possibility is to match mechanically to the glass rather than the silicon. From the table above, borosilicate glass has a speed of sound of 5.3km/s, silicon has a speed of sound of 8.0km/s, and pyrolytic carbon has the very low speed of sound of 1.5km/s. If the silicon is thinned from 100 microns to 95.8 microns, and the remaining 4.2 microns is replaced with pyrolytic carbon, then the average speed of sound vertically through the two layers is the same 5.3km/s as the glass. The "elastic impedance" is proportional to the energy stored by propagating sound, and sound waves will reflect at impedance discontinuities. The impedance of borosilicate glass is 11.9 $Kg/mm^2 s$, silicon is 18.7 $Kg/mm^2 s$. The composite material is 12.4 $Kg/mm^2 s$, a much better match to glass than pure silicon, so standing waves at the interface are less likely. Finally, pyrolytic carbon is an excellent thermal conductor. The composite structure has 50% better thermal conductivity than silicon alone.

There may still be issues with the Si/pyroC composite material; it may help or hurt adhesion between server sats. black body radiation, trapped charge, etc. The stacked material is acoustically dispersive for high frequencies, and this may help with shatter resistance and handling. The carbon help may damp out vibrations faster. The final server-sat stack will consist of many layers of materials, and they will be evaluated empirically to find the best mix.

Curling can occur if the front and back sides of a server-sat (especially the solar cell) have different CTEs, and the server-sat undergoes repeated thermal cyclings. There is nothing inherent in a server-sat that establishes "flat" - it will flex until tensions and compressions are minimized. A slightly curved server-sat is not a severe operational problem. If the edges are turned up a few degrees, that will reduce collected solar energy very little. The main problem is the effect on the phasing of the radios. Significant curling will change the spacing of the radios at opposite sides of the curl, and lift them above the plane of the radios at the center of the curl. Without some means of determining precisely how much curl is there, the radios may be in incorrect phases.

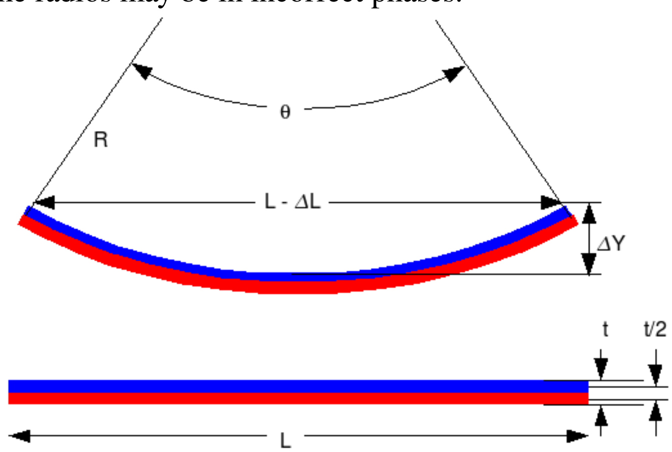

If $\beta = \Delta CTE \times \Delta T$, and the materials have equal thickness and Young's moduli (the worst case for curl), then the center-line of the top (blue) material will be approximately $L(1- \beta/2)$ and the bottom (red) material will be $L(1+ \beta/2)$. These correspond to an arc with a radius $R=1/(2\beta)$ and an angle $\theta = L/R = 2\beta L / t$. This leads to $\Delta L = L=2R(\sin(\theta/2)$ and $\Delta Y = R(1-\cos(\theta/2))$. Using the Taylor series expansions for sine and cosine, these can be approximated as:

$$\Delta L \approx (1/6) \times (\beta/t)^2 \times L^3$$

$$\Delta Y \approx (1/4) \times (\beta/t) \times L^2$$

Here are some curls for various temperature and CTE changes and sizes:

L (mm)	t (mm)	ΔCTE	ΔT	β	ΔL (mm)	ΔY (mm)
300	0.10	1.00E-6	100	1.00E-4	4.50	22.50
300	0.10	5.00E-7	50	2.50E-5	0.28	5.63
300	0.10	2.00E-7	100	2.00E-5	0.18	4.50
200	0.10	1.00E-6	100	1.00E-4	1.33	10.00
200	0.10	5.00E-7	50	2.50E-5	0.08	2.50
200	0.10	2.00E-7	100	2.00E-5	0.05	2.00
150	0.05	1.00E-6	100	1.00E-4	2.25	11.25
150	0.05	5.00E-7	50	2.50E-5	0.14	2.81
150	0.05	2.00E-7	100	2.00E-5	0.09	2.25

Factories

Much of the manufacturing for Server Sky can happen in Washington County, Oregon, around Hillsboro.

The main component by weight is the solar cell. The Solarworld factory in Hillsboro [SOLA] is the largest solar cell manufacturer in the US - certainly the most highly automated. The solar cells in their illustrations appear to be 100 millimeters diameter, but perhaps they can make larger ones.

The most complicated components are the microprocessors. Some version of the Intel Atom [ATOM] may be suitable. For a server-sat, it is preferable to use a fast, deep sub-micron, 1V processor with heavy doping (less sensitive to radiation damage) and an epi substrate, but preferably trench isolated SOI. AMD processors are all trench isolated SOI. Intel's Penryn process, with thick Hafnium oxide gates and work-function controlled gate metalization, will also be more radiation resistant. Over time, most process improvements desirable for high performance processors will also be desirable for radiation hardness.

The most radiation sensitive components are likely to be the flash memory. These incorporate error correction, but software error correction and frequent rewrites may be necessary to correct for radiation induced charges.

The gallium arsenide radios may be made by TriQuint [TRIQ].

The electro-chromic material will be pretty simple. It must survive freezing, and should be divided into separately-addressed cm-sized cells so that meteorite punctures will not disable a whole thruster. They will probably be constructed from a 1 micron layer of electro-chromic material between two pieces of indium oxide coated 30 micron thick glass (which is commercially available). This is standard technology. Similar LCD material is currently manufactured by Sharp in Vancouver, Washington.

Washington county companies such as D.W. Fritz [DWFR] build wafer handling equipment .

Radiation

Server Sky

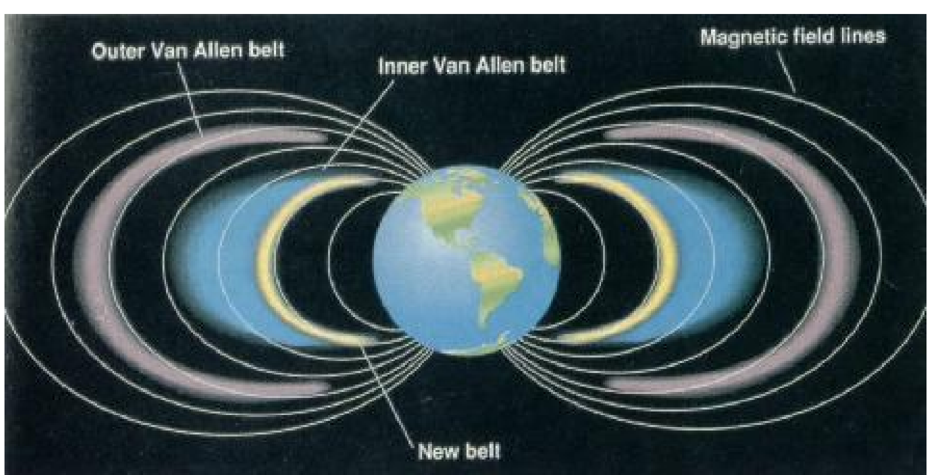

The M288 orbit (4 hour sidereal, 5 orbits per day relative to the earth) is located at a radius of 12789 km, about 6400 kilometers altitude. That places it between the inner and outer van Allen belts. This is a relatively high radiation environment compared to low earth orbit, but lower than the center of the belts. Recent advances in semiconductor technology permit unshielded electronics to operate in this region, and Server Sky takes advantage of that work. Traditional satellites are damaged by this much radiation; consequently, there are few satellites operating in these orbits.

Ionizing radiation does five nasty things to semiconductors:

- Latchup
- SEU - single event upsets, bit flipping
- Oxide charging
- Flash memory errors
- Crystal lattice degradation

Server sky is designed to resist these effects. The space radiation environment is well characterized. Computer models exist, such as AF-GEOSPACE [AFGE] Radiation effects can be empirically tested in ground laboratories, at high dose rates, then extrapolated to the effects of years of space radiation exposure.

Latchup

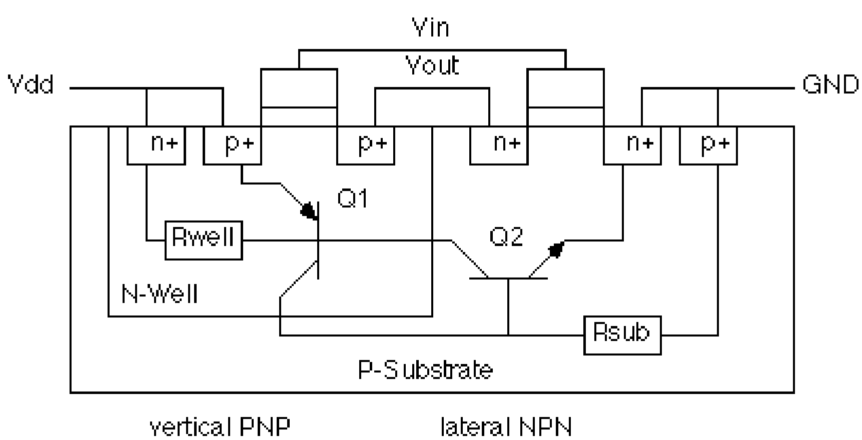

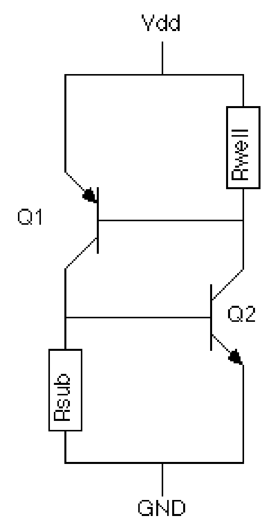

The NMOS and PMOS transistors in a CMOS integrated circuit are often isolated from each other by diode junctions. The first drawing above shows a PFET in an N well next to an NFET in a P well. A P source, N well, P well, and N source make a PNPN junction stack, as shown in the second drawing. This is much like a device called a Silicon Controlled Rectifier or SCR, which is a commonly used latching electrical switch. When the currents are high enough in an SCR, a positive feedback loop is established, latching the device in a conductive state. In this case, that can form a high current conducting path between the power supply and ground. This heats up the silicon, and can destroy it.

An ionizing radiation particle can create a temporary high current path that can activate an SCR and cause latch-up. This is a problem in electronics exposed to radiation. It is usually mitigated by special design for latch-up resistance, and heavy shielding for the electronics. Server Sky cannot afford the shielding weight, or the special design.

The SCR path cannot be activated if the voltage between the supply and ground is less than a diode turn-on voltage (typically more than 0.7V). The electronics in a Server-Sat are powered by a single-junction solar cell, which is a large forward-biased silicon diode. It cannot produce more than about 0.6V. A **server-sat cannot latch up!**

In some special circumstances, circuits called "voltage multipliers" may be used to construct higher voltages for special needs. However, the voltage multipliers will be low current and easily switched off, so damaging high current latch-up will not be possible in these circuits, either.

Single Event Upsets

The charge deposited by an ionizing particle can temporarily overwhelm a logic gate, or change the state of a register bit. This can cause incorrect computations. Often the cost of such an error is small, for example when computing the I and Q signal of a software defined radio - a calculation error will just create a little extra noise. The cost is unacceptable when determining CPU state, performing calculations, and such. While computation errors can be detected and corrected with redundancy, either by duplicating hardware or by repeating calculations, redundancy is lowers performance and computational efficiency.

One promising approach is to use the Razor error correction technology [RAZ] being developed by the University of Michigan, MIT, and Intel. Digital integrated circuits are typically designed for high "noise margin", with extra power and voltage swing added, and clock rates reduced, to reduce the chances of a logic failure to infinitesimal probabilities. However, this is costly and reduces performance. RAZOR reduces the noise margin, greatly improving the performance, but at the cost of frequent errors. RAZOR adds circuitry to detect these errors, and repeat the calculations when they occur. The system can be tuned so that the performance improvement greatly exceeds the cost of the extra calculations, doubling overall performance. RAZOR technology will be common in logic chips and processors in a few years.

Radiation-initiated single event upsets are just another kind of error, which can be detected and corrected by RAZOR-like technology. Since these SEU events may occur in the error correction logic itself, there will need to be additional redundant logic in the correction circuitry, but that will still be a net win over adding redundancy to the entire chip.

Gate Oxide Charging

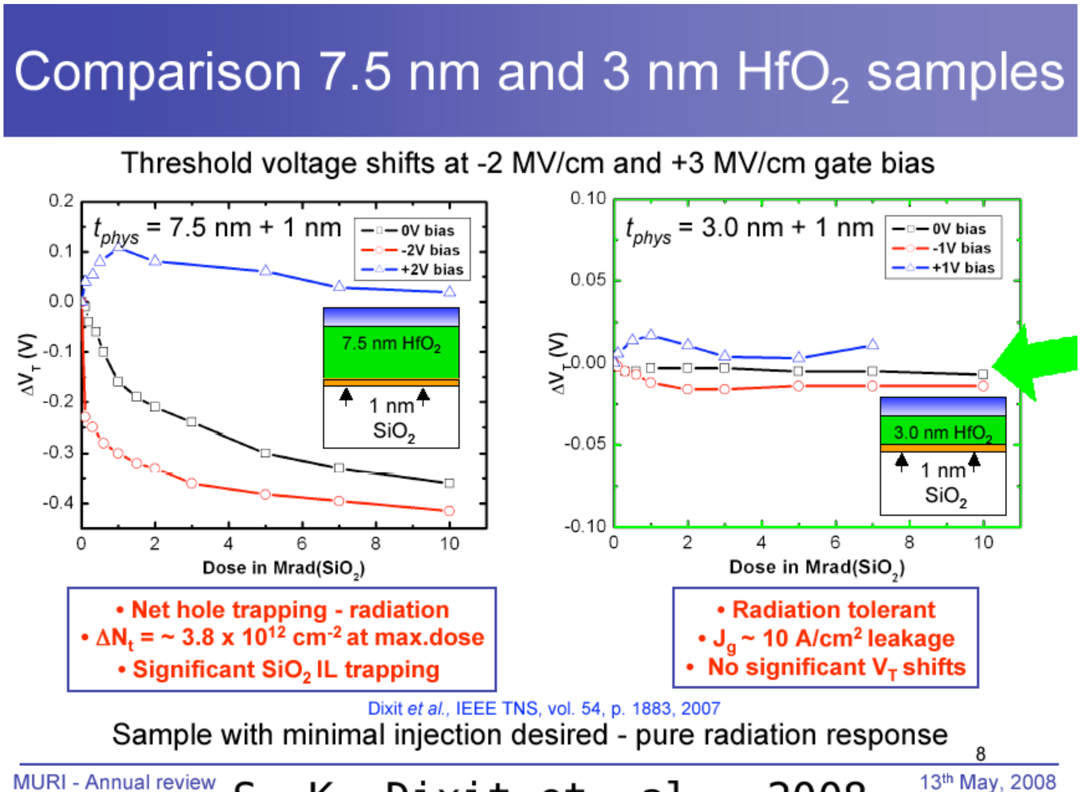

S. K. Dixit et. al., 2008
Vanderbilt University, Intel Hillsboro

Silicon dioxide develops a positive charge when irradiated. An ionizing particle passes through, and generates hole-electron pairs. The electrons are highly mobile, and diffuse or drift out, while the holes get trapped, and leave a positive charge. Hafnium oxide develops a negative charge, trapping electrons. Recent work by Dixit shows that a stack of both shows promise as a rad-hard gate oxide, withstanding 10Mrad from a Cobalt 60 source with minimal shifts.

Space Junk and Debris

Collisions are a problem, especially for arrays like server sky which will put huge numbers of assets into carefully defined and controlled orbits.

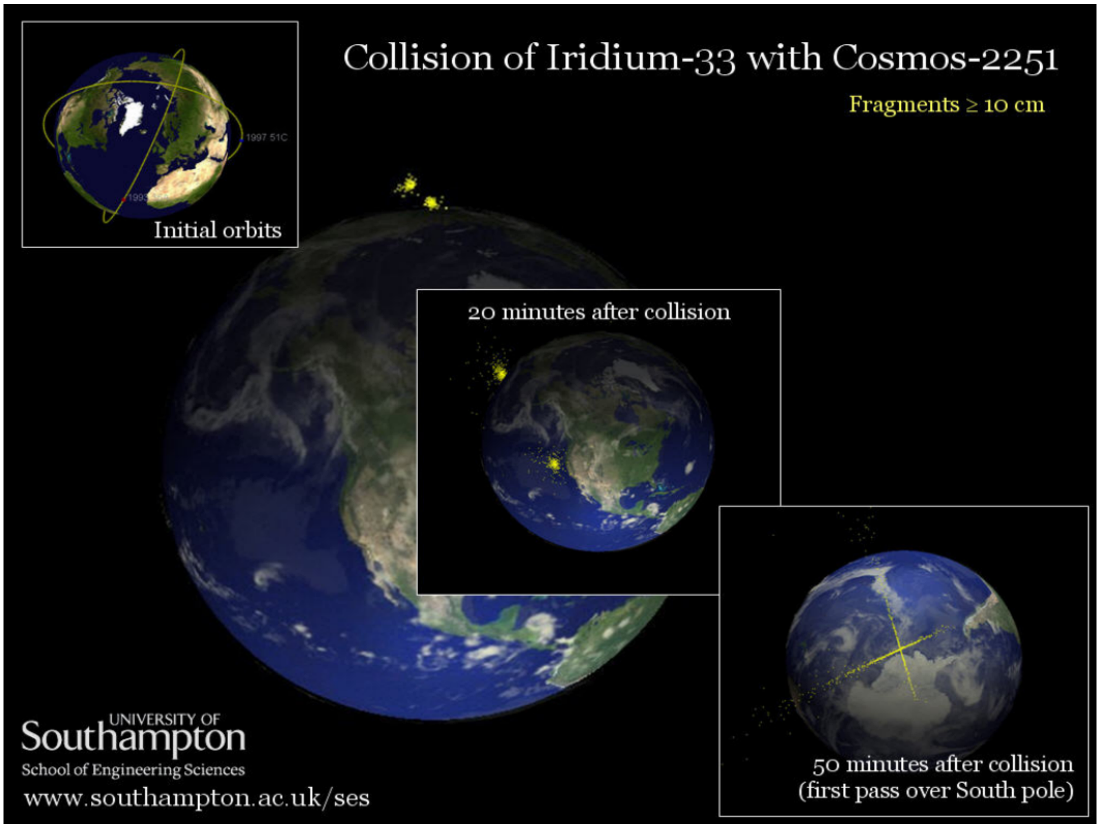

The above picture illustrates the 2009 February collision between Iridium-33 and the abandoned Cosmos-2251 . The collision destroyed a $250M asset, and created hundreds of large pieces of debris. Server Sky cannot be permitted to come close to such expensive assets as the Iridium communication satellites. The consequences of a collision are too costly.

On the other hand, server-sats are cheap, and destroying one in a collision is unfortunate but not especially expensive. The main cost of a server-sat collision with space junk is more space junk, which can damage other server-sats. More robust objects (like traditional big-iron satellites) encountering server-sat fragments will probably lose a solar cell or two on the skin, but won't be destroyed. Still, it is bad manners to add to the debris problem.

Fortunately, server-sats can be maneuvered out of the way of predicted collisions, if the predictions are accurate enough. Since they maneuver by light pressure, they have an unlimited supply of very low thrust - given a few hours, they can change orbit sufficiently to stay more than a kilometer away from any tracked impactor.

Server sky will be deployed in orbits higher than most space junk. The version 0.2 design deploys server-sats in m288 orbits, with semi-major axes of 12781 kilometers and average altitudes of 6411 kilometers. The vast majority of space debris is in lower orbits - it requires high launch velocities to even reach those altitudes.

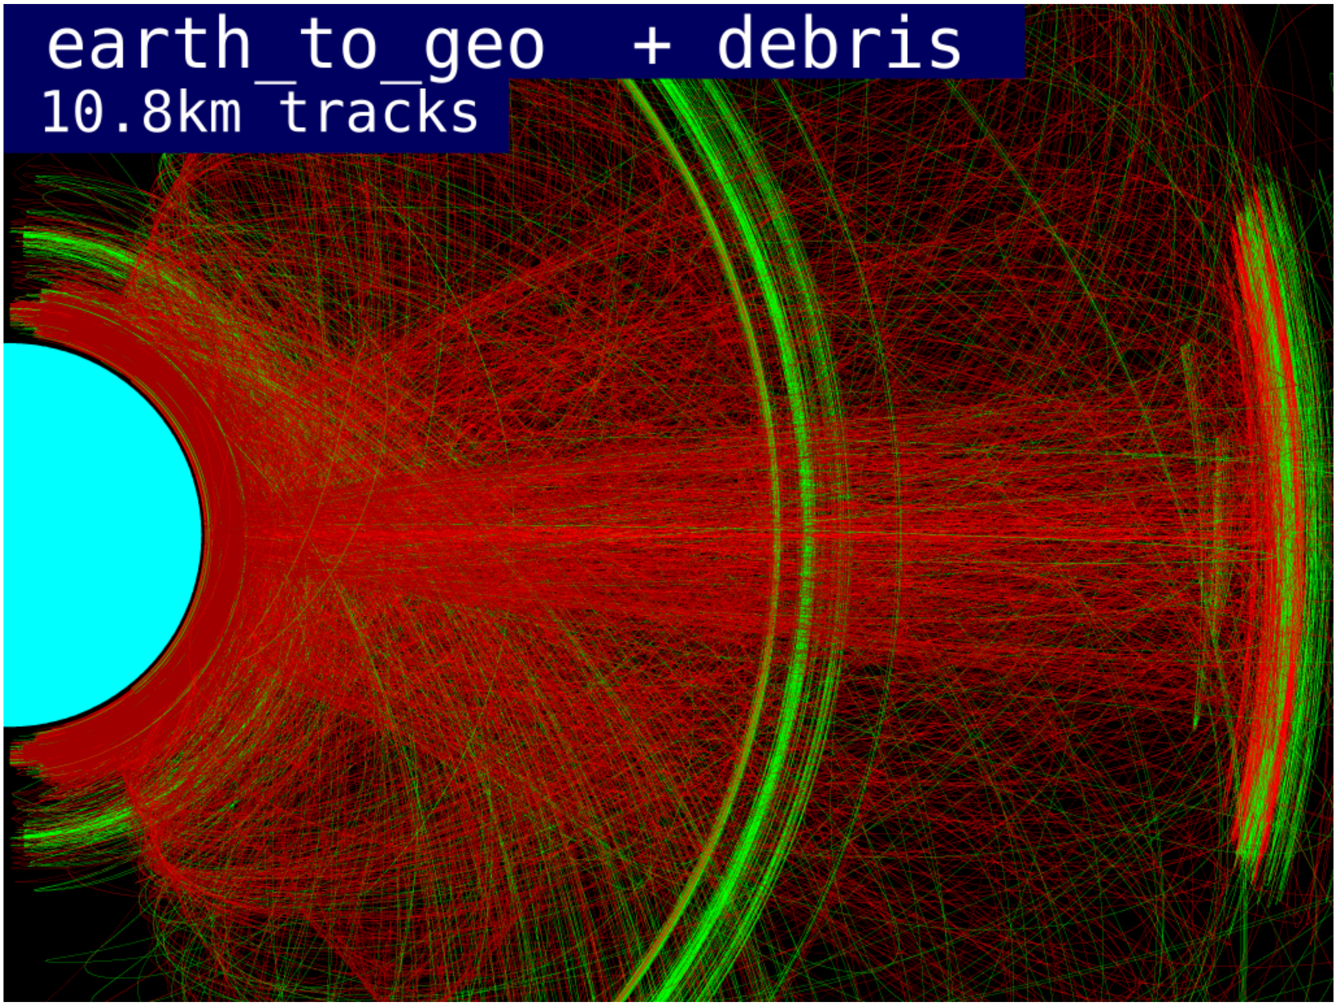

The above plot shows all the objects tracked by NORAD in orbits up to geosynchronous altitude. This is an "HV" plot, with the horizontal H component representing the radius in the equatorial plane, and the vertical V component representing the distance above the plane. All orbits will look like an ellipse or banana shape plotted in this way. Red is debris, darker red is spent rocket bodies, and green is useful satellite assets. At any given time, the objects are in one place on these orbital tracks.

The thick cloud near the earth is LEO or Low Earth Orbit. These are the easiest orbits to reach, and the vast majority of useful satellite assets, and most of the space junk, is located here. The stripes on the right of the drawing are satellites in geosynchronous orbits. These orbits are well managed, with satellites actively positioning themselves, and spent satellites and rocket bodies in higher and lower orbits.

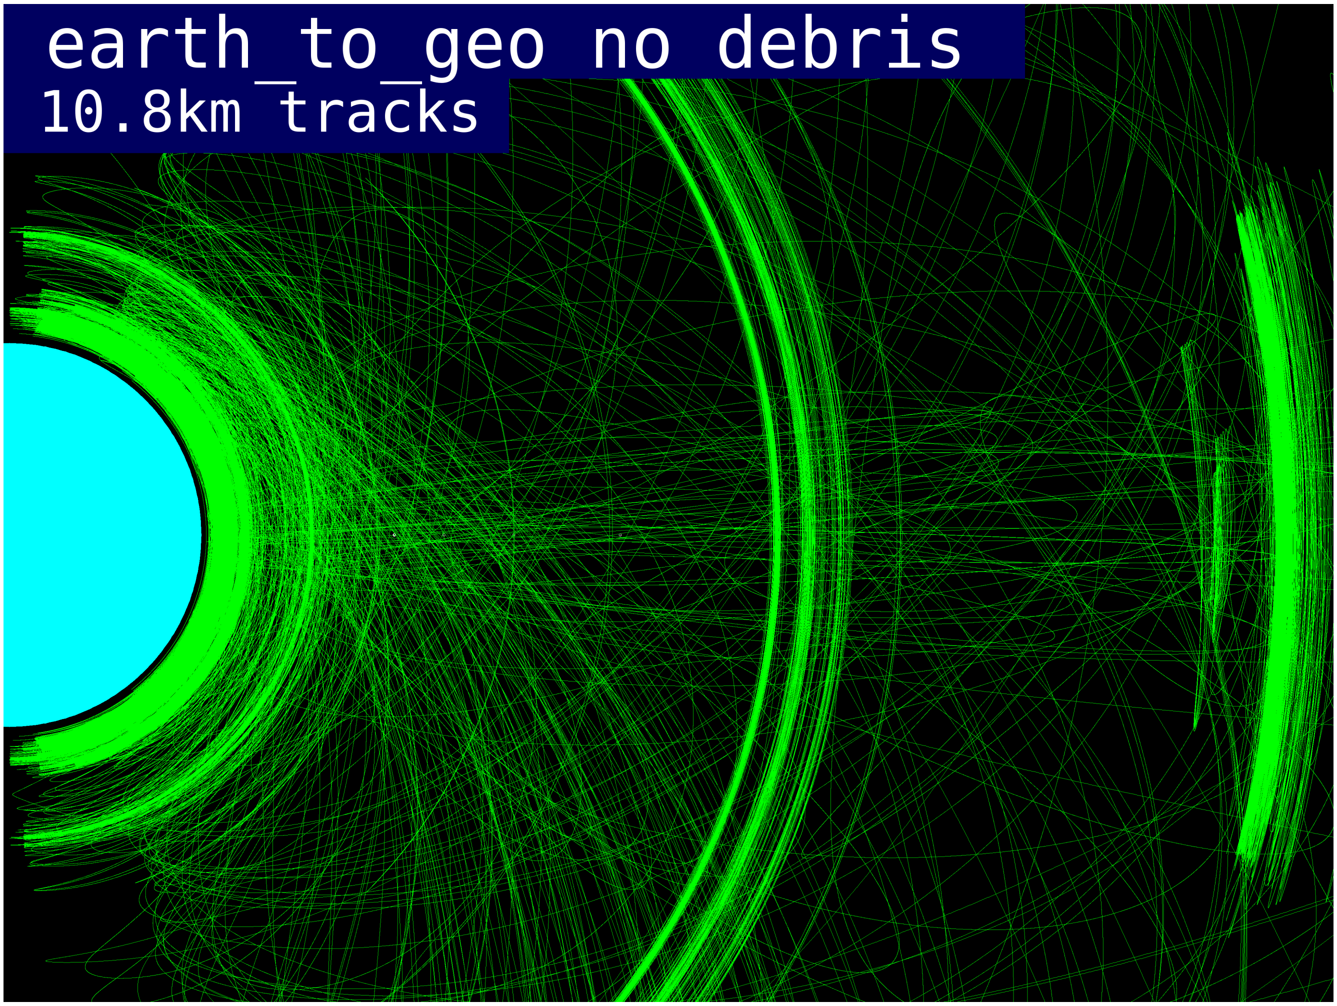

Above is the same plot to GEO with only "intentional" satellites shown; far fewer objects than the debris plots. Satellites have been accumulating for 50 years, and only the satellites with low perigees ever decay and fall to earth. So, most of these satellites have been in orbit long after their useful life, and many are out of thruster fuel or completely dead Only a fraction of the "green" satellites are valuable assets.

In between GEO and low earth orbit is "MEO", or Medium Earth Orbit. The first interesting band of MEO is the green band extending above and below the poles, a bit further out than LEO. This region contains the Iridium and Globestar communication satellite clusters - hundreds of satellites that form a communication mesh for telephone and other low latency communication. Sadly for the operators of these systems, fiber optic communication got a whole lot better, and in most parts of the world that is a cheaper way to move telephone and data globally. The operators of Iridium and Globestar went bankrupt.

The bright green bands running through the center of the drawing are the GPS/Navstar and GLOSNAS navigation and positioning satellites.

The distances represented by this plot is vast - 44000 kilometers wide, 33000 kilometers high, more than 10 times the land surface of the earth. The volume of space represented is enormous, 2E14 cubic kilometers, or 200 trillion cubic kilometers, or 400 thousand times the volume of atmosphere in which we fly all the airplanes in the world. The tracked items are small from centimeter scale (debris) to 100 meter scale (international space station). So an "accurate" drawing would be completely blank.

There are far more airplanes airborne at any given time than satellites, and most satellites are much smaller than most airplanes. So why are there any collisions at all? There are occasional rare collisions because the orbiting bodies move continuously at high speed, the orbits differ in period and relative inclination, and after satellites fail or expend all their thruster fuel, they cannot avoid predicted collisions. Airplanes going in different directions are assigned to different altitudes, so they should never intersect with high closing velocities. Airplanes are actively tracked with radar, and change course to avoid other airplanes. Satellites follow no such discipline, but they could.

NORAD tracking information is made with ground-based radars in a very few locations around the earth - most of Low Earth orbit is unobserved. Atmospheric drag on the lowest objects can significantly alter the orbits. With current capabilities, it is very difficult to predict future satellite orbits better than a kilometer or so. Since thruster fuel is limited, avoiding every low-probability collision event is a futile exercise. When we develop the ability to track objects to much higher precision, we will be able to accurately predict future collisions, and make small thruster changes to avoid them.

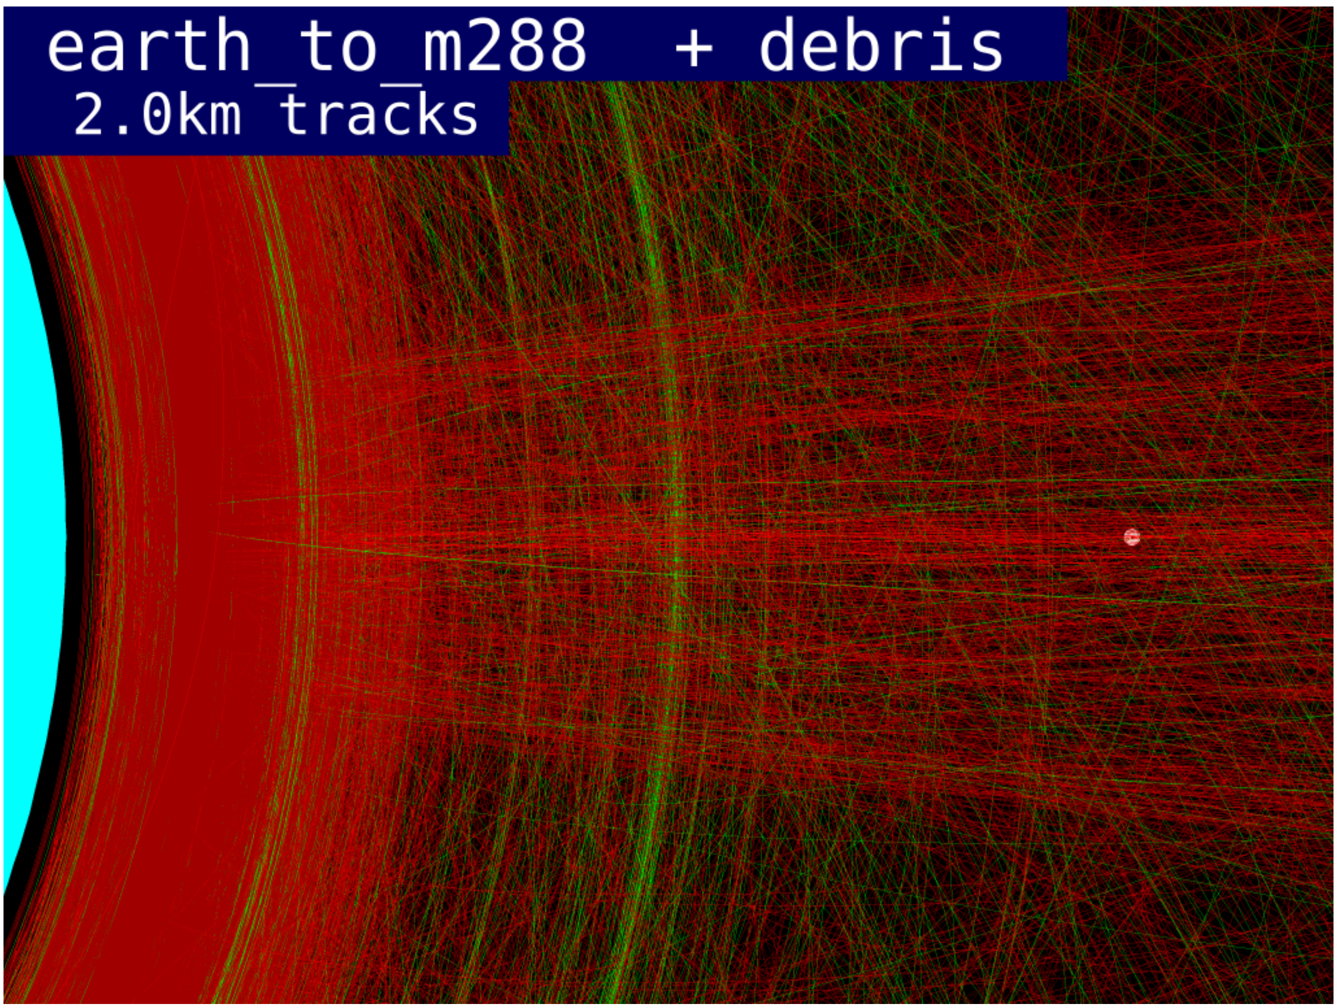

Server sky version 0.2 will be located in m288, about one earth radius in altitude. The 100 kilometer dot to the right of the drawing above represents the orbits used by server sky. There is still some debris and a few useful assets at that altitude, but it is not nearly as crowded as low earth orbit.

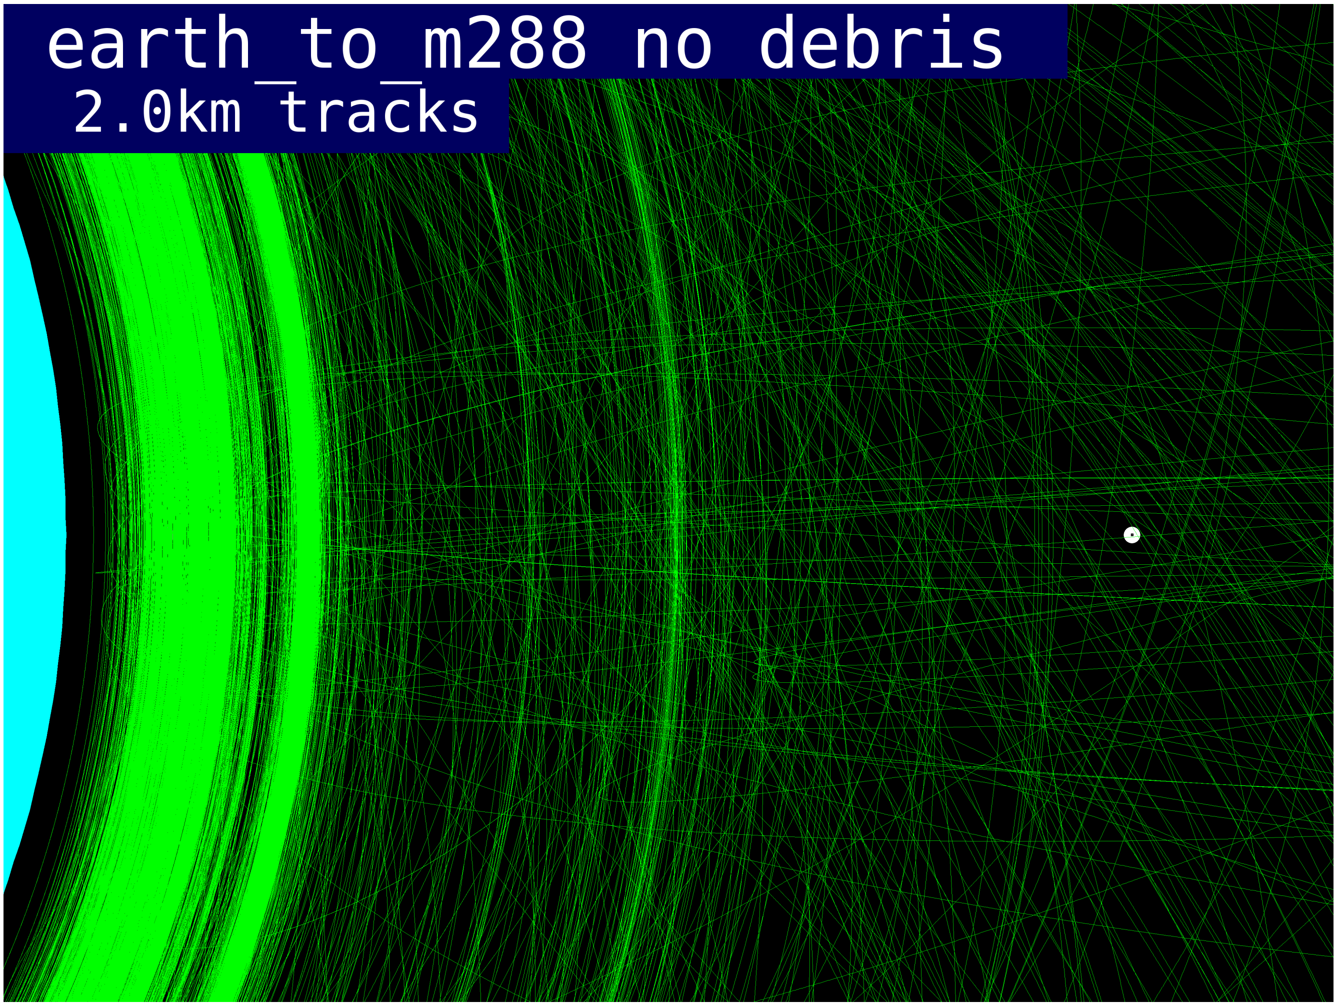

Above is the same plot without the debris. Only a few assets ever get close to the m288 orbit at 6411 kilometers altitude. Except for the bands of navigation and communication satellites above, the density of satellites tends to drop off exponentially with altitude. Keep in mind that these plots compress in the circumferential direction - the collision density is half as large, measured in satellites per cubic kilometer. Collisions are also proportional to closing speeds, and orbital velocities are approximately proportional to the square root of the radius. Lastly, most orbits that extend this high are not subject to unpredictable drag from the atmosphere, and are high enough to be frequently above the horizon for ground-based tracking radar. Their orbits can be precisely characterized. The collision problem is not absent, but greatly reduced at altitudes above 1000km .

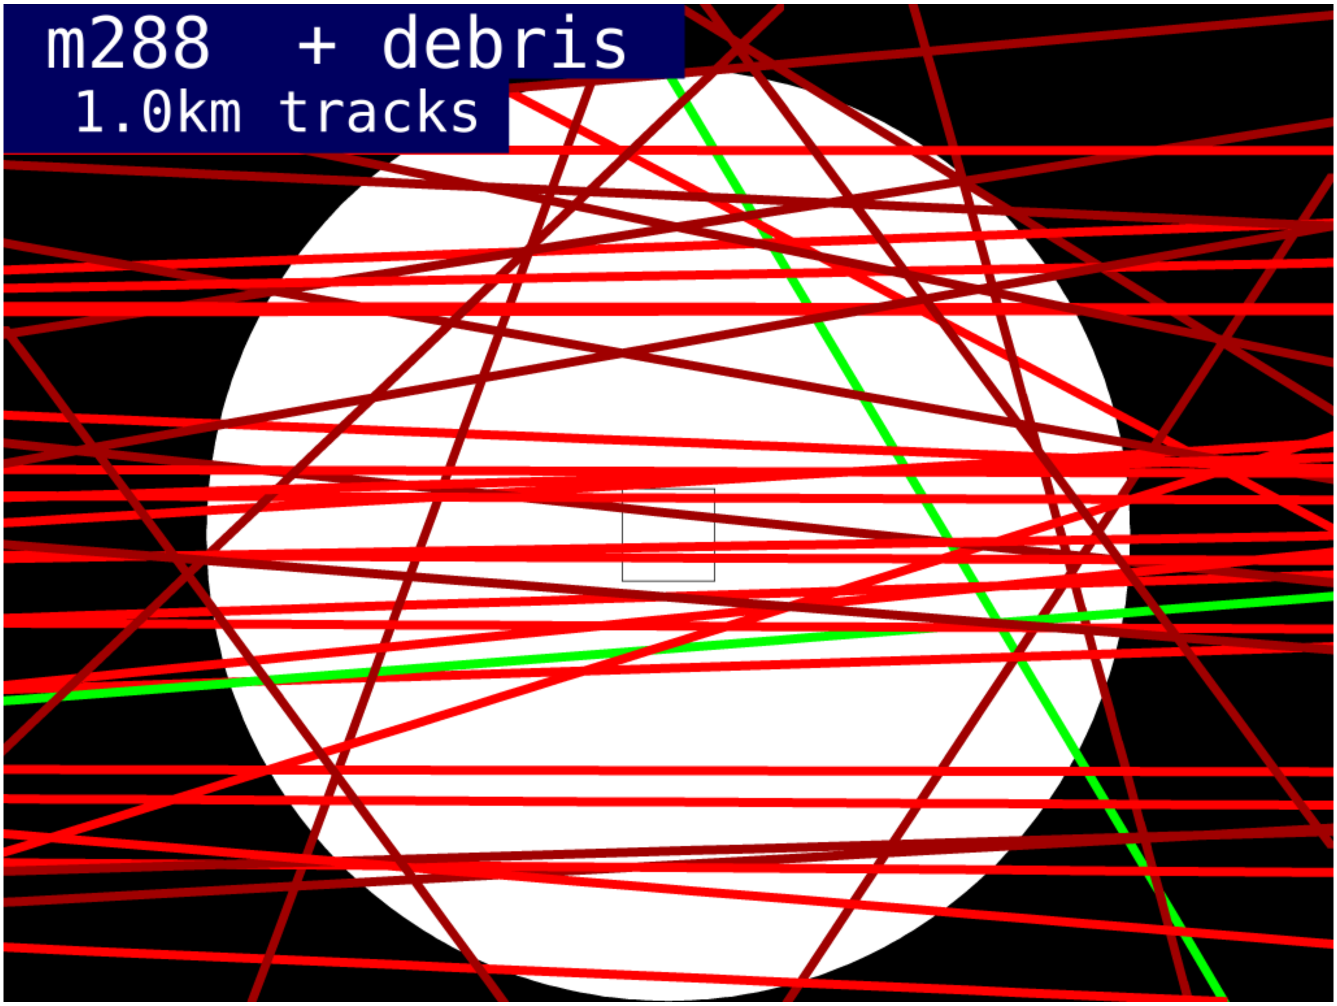

Above is a plot of just the m288 region. The square in the center of the plot is 10 kilometers on a side, about as big as the city of San Francisco. If the Server Sky orbits are threaded through the white spaces, they will *never* encounter one of the NORAD tracked space objects, unless its orbit shifts slowly over time. None of these objects are synchronous with the server-sky constellations, but until m288 becomes crowded with billions of server-sats, the constellations can be moved a few kilometers up or down orbit to give these potential colliders a wide berth.

NORAD does not track everything, especially small sub-centimeter objects that are hard to detect with ground-based radar. Orbiting radars - especially if there are many of them - can do a better job, especially if they are using very short wavelengths and large aperture antennas. However, until that detailed mapping is available, centimeter objects that might destroy a solar cell or two on a large satellite might shatter a server-sat. Perhaps server-sats can be designed to not shatter when a sub-centimeter collider passes through them - if all it hits is a segment of the solar cell, the server-sat can keep functioning with slightly less power.

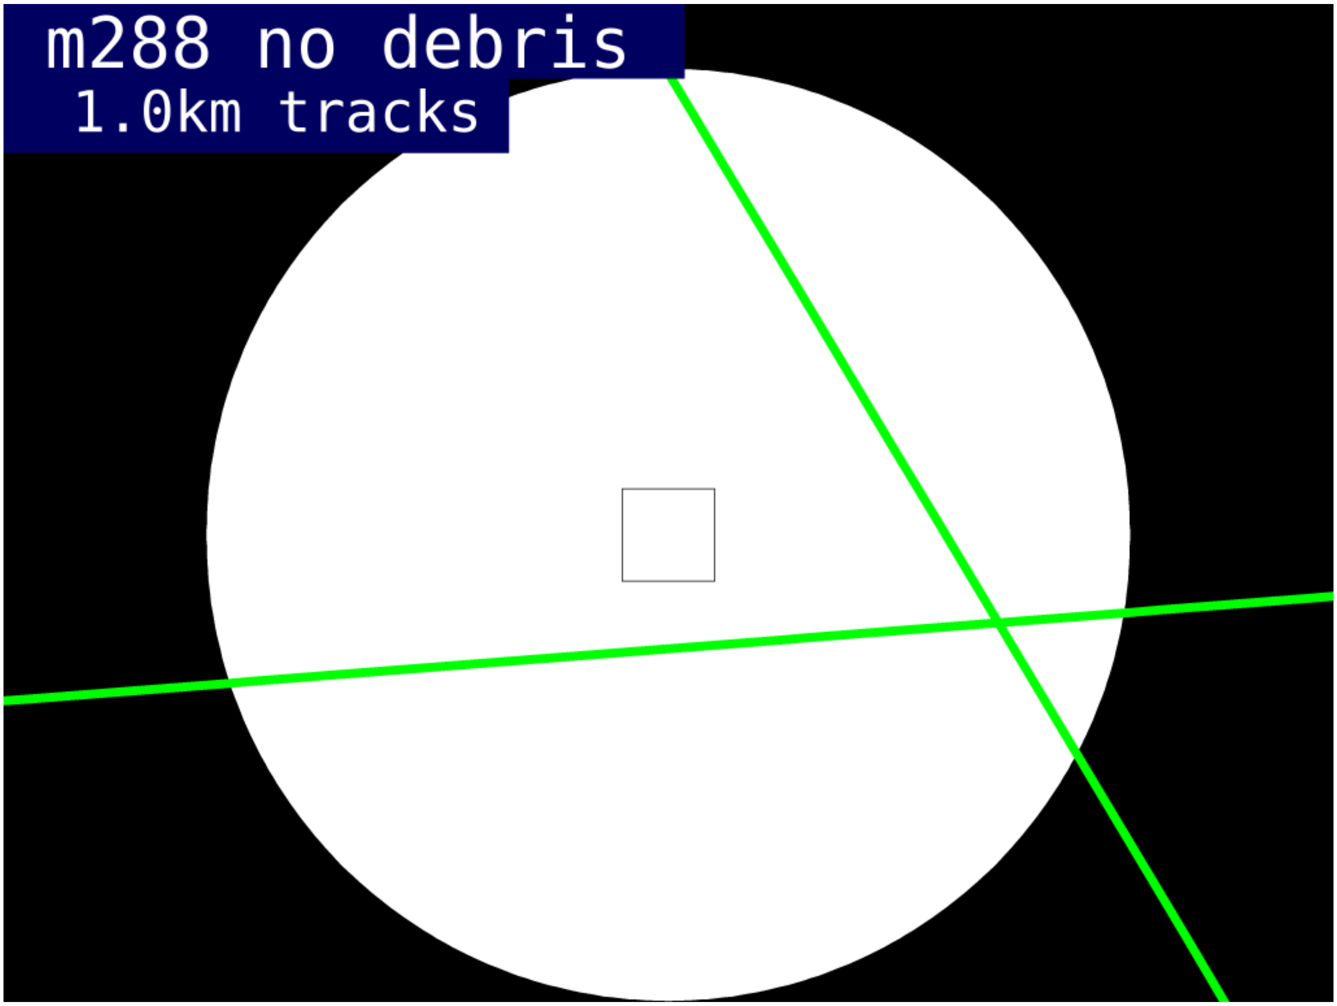

Above is the same m288 plot without the debris. Only two assets cross the server sky orbit at this time; MOLNIYA 3-3 and BSAT 2B. The Molniya was launched in 1975 and may no longer be in use. The BSAT is a Japanese communication satellite that failed to reach geosynchronous orbit. So even these are not valuable assets. If the debris and failed satellites are removed, the orbits should be completely clear and usable for Server Sky. Until then, accurate tracking and avoidance should free up most of the orbit.

Surface charging

Most of the area of a server-sat will be solar cells, with low voltages across the surface. If the electrochromic thrusters have glass covers with Indium oxide conductors on both sides, the outer conductors are grounded relative to the rest of the server-sat. As all surfaces will be conductive, and at voltages less than one or two volts, there is no opportunity for arcing across surfaces, which sometimes happens to geosynchronous satellites in the solar wind (see [SMAD] page 212-214).

Night Side Temperature

When the server-sat orbit takes it to the night side of the earth, for about 1/6th of its orbit, it will only be heated by the thermal radiation from the earth, filling 0.5*(1 - cos(30°)) or 6.7% of the sky with a black body temperature of 250K. At equilibrium, the temperature of the server-sat will be 250K times the fourth root of 0.067, or around 127K (-146C).

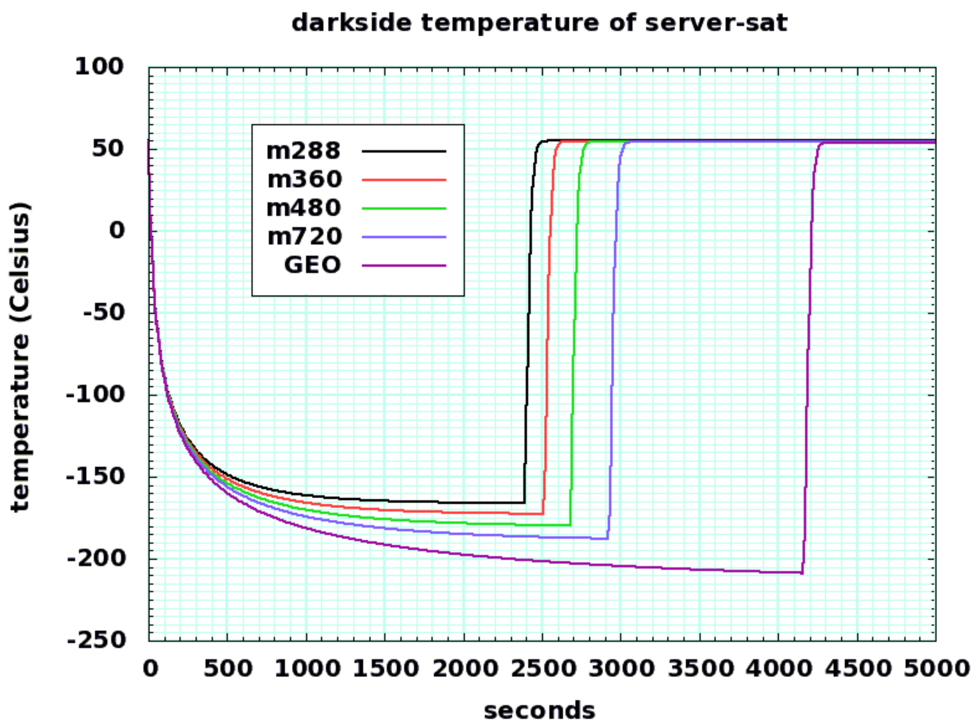

Here is a plot of how the temperature varies as a server-sat passes behind the earth. The dark time corresponds to the spring and autumn equinoxes; in winter and summer, the shadow time will be a little smaller for m288 orbits. In GEO and higher orbits, there will be no shadow at all in winter and summer, though there will still be shadows in spring and fall.

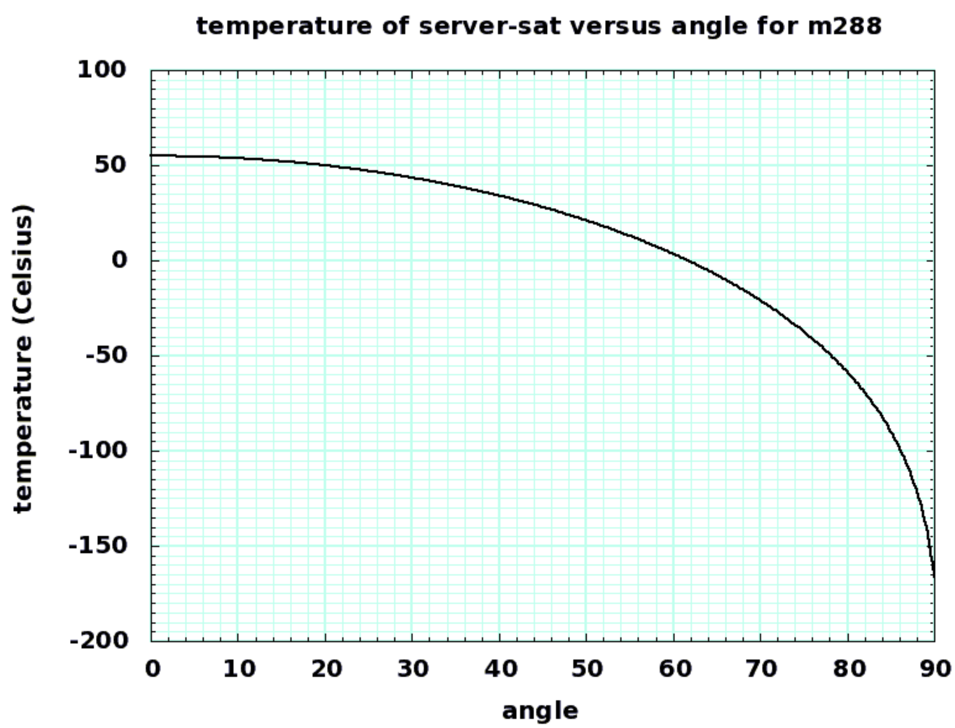

Turning the server-sat means it intercepts less solar energy. This will cause it to cool down.

Possibilities

This is long term speculative stuff. Mostly it shows that we can bypass most conceivable limits to growth if we move most of our power production, communications, computation, and industrial processes into space and away from the biosphere.

High Orbit Arrays

For the first decade or two, server sky will coexist with traditional geosynchronous satellites, but as the traditional satellites age and become a very small part of the total communication bandwidth around the earth, they will likely be replaced by more server-sats.

Terascale Arrays and Beam Power

For non-realtime calculation, big compute jobs like weather prediction or animation rendering, a much larger latency is tolerable. Server-sats in the Earth-Moon Lagrange positions will be 60 times further away from earth, so under similar conditions as before, they will produce 1/3600 of the worst case night-time illumination as a server-sat in m288. There is room for trillions of server-sats in these Lagrange positions, with round trip ping times of 2.5 seconds.

This is far more than is needed to provide foreseeable computation and communication needs, so many of the later generations of server-sats may become "compute-light" and "transmit-heavy", beaming the power as microwaves to rectenna arrays on the ground, producing power for the electrical grid. Because the microwave beams are steerable, they can move from peak load center to peak load center as needed, reducing long-lines requirements. They can even be steered in circles around 6 rectenna grids, generating 3 phase AC power. This is an old idea - solar power satellites - but arrays of server-sats are much lower mass, cheaper, and easier to deploy than large rigid systems of solar cells, structure, and antennas.

However, high density microwave beams are not healthy. They can be stopped by a thin layer of metal, but birds are not shielded. Unless birds can be reliably kept away from the rectennas, the rectennas should only be placed where birds aren't. Perhaps the best place for rectennas is over deep ocean, far from land and far from the paths of feeding and migratory birds. A few centimeters of ocean water will stop the power that leaks through the rectenna, so sea life is safe from the leakage while still getting sunlight.

With good ground-based telescopes and pattern recognition, and huge orbiting arrays with very small ground spots, it should be possible to create nulls in the ground pattern where individual or flocks of birds are observed. With 2.5 seconds of latency in beam steering from L5, the nulls must include not just where the birds are, but where they can get to in the next 3 seconds. It may be possible to constrain the bird's flight path with small robot airplanes, noises, and other stimuli so that the behavior of the birds is a little more predictable and constrained.

With these and many other constraints, it should be possible to supply most of the Earth's energy needs from huge arrays of server-sats in space.

Lunar and Asteroid Materials

Most of the mass of a server-sat is silicon, glass, and aluminum, which are also most of the mass of rock, including lunar rock. It may be a long time before we can manufacture solar cells off the planet, and much longer before we could manufacture deep sub-micron integrated circuits. But the materials on the moon are in a lower gravity well, and there are few ecological risks to using large amounts of lunar material to manufacture solar cells and glass substrates, and launching them with electromagnetic launchers. A cubic meter of lunar regolith could be used to manufacture perhaps half a million 3 gram server-sat "chassis", which could be mated to earth-manufactured integrated circuits in an automated facility in orbit. A cubic kilometer of regolith could manufacture half a trillion server-sats. Lifting those server-sats off the moon and placing them in an m288 orbit would require about 1e19 joules (depending on how they were captured), which is about as much energy as the finished product would produce in half an hour. Astoundingly rapid investment rates of return are possible with lunar materials.

Deep Space Arrays

There is a lot of room in the solar system. Outside the orbit of the earth, most of the light is dumped into interstellar space. Server-sats orbiting between Earth (1.5E11 meters from the sun) and Mars (2.3E11 meters) could capture much of the light of the sun. If there were enough of them, it would increase the apparent infrared temperature of the sky, which would in turn increase the temperature of the Earth. If the earth temperature increase was limited to 1C, then the effective sky temperature could increase from 2.7K to 100K. If the server-sats were at 1.9E11 meters distance from the sun, receiving 800 watts per square meter and at an equilibrium temperature of 270K, then they could cover about 2% of the sky. That intercepts about 7E24 watts of light, and might generate about 1E24 watts of usable electric power for computation and industrial purposes. Arrays near Jupiter would receive far less light, but would pose no significant infrared problem.

Low cost launch

The Launch Loop [LOOP] is an electrically powered earth-to-high orbit launch system. The main construction and operating cost of a launch loop is electricity. At 10 cents per kilowatt hour, and a quick payback of capital, a launch loop can put a kilogram into orbit for about $5, and a small launch loop can launch 80 tons into high orbit per hour.

Assuming extra mass for the satellite bus and the apogee insertion motor, the cost of orbiting a 7 gram, 2 watts-to-ground-collector server sat will be on the order of 12 cents. If that 2 watts can be collected for another 10 cents of rectenna infrastructure, and the mechanism that does so lasts 20,000 hours, that is 100 kilowatt hours per dollar invested. This drops the cost of further launches. Thinning the server sats down to 1 gram will save more. In time, the cost can drop still more by building apogee capture systems such as rotating tethers (with some payloads sent around the moon to add momentum back to the tether system).

The result will never be free space launch, or "power too cheap to meter", but it can result in very low cost space launch and electric energy on the earth - 50 cents per kilogram, and a 50 cents per megawatthour, may be possible someday.

Conclusion

Server-sky is speculation. There are many unsolved problems, and more will surely crop up during implementation. Fortunately, the problems encountered so far have shown signs of solution. With enough imaginative contributors, other problems and their solutions will emerge, often from unlikely places elsewhere in the world.

Server-sky may be the near-term commercial application that will pay for large scale space development, leading to the permanent expansion of earth-life into space. It may also be how we save earth-life from destruction, by moving large scale computing, and eventually power generation, into space.

This paper is a plea for your participation. If Server Sky is as important as it seems, the visionaries that design, build, and operate Server Sky will be remembered by history as the heros that saved mankind. The idea is in its early stages - you don't have to be a rocket scientist to contribute important new ideas (we can use those, too!), but you must be willing to critically analyze you own ideas - don't just make noise and pompous announcements. Do the math. Do the research. Read. Observe. Write. Code. Draw. Animate. Share. Much of the work will be difficult, but that is why you have a brain. Use it.

Server sky will be developed as open technology. Major corporations will participate, and many will earn vast fortunes providing products and services involving server-sky. But the idea will affect the entire world, and the world must be engaged in doing it right. If we do this in secret, we will face angry saboteurs, social as well as physical, when we go public. If we do this together, we will succeed together.

References

This paper is incomplete. Much more information, and more recent versions of this paper, can be found on the wiki at **http://server-sky.com/** . The website is a wiki, please feel free to correct or enhance it. Rants will be moved, or removed.

[AFGE] http://www.kirtland.af.mil/library/factsheets/factsheet.asp?id=7899 AF-GEOSPACE is a collection of programs that model the radiation environment near earth

[ALIX] http://wiki.keithl.com/index.cgi?SL5Alix The 4 watt ALIX board from PC Engines, using an AMD Geode X86 processor, producing 900 bogo-MIPS performance.

[AREA] http://www.osti.gov/bridge/product.biblio.jsp?osti_id=764362 An estimate of road and roof area in Sacramento; multiply by the population ratio of 230 for the US. WAG.

[ATOM] http://www.intel.com/products/processor/atom/ Intel Atom low power processor

[BYUC] http://www.ece.byu.edu/cleanroom/CTE_materials.phtml Useful numbers for the coefficient of thermal expansion of various electronics-related materials. Note that crystalline materials are often isotropic.

[DATA] http://www.energystar.gov/index.cfm?c=prod_development.server_efficiency_study E.P.A. estimate of data center power usage.

[DEBR] http://apollo.cnuce.cnr.it/rossi/publications/iau/node2.html A 1994 estimate of space debris. There is more now, but it follows the same pattern.

[DIXI] http://www.isde.vanderbilt.edu/content/muri_2008/dixit_muri2008.pdf Sriram Dixit et. al. at Vanderbilt University. Recent work on HfO/SiO2 stacked gates and radiation resistance.

[DWFR] http://www.dwfritz.com/ D.W. Fritz makes semiconductor wafer handling equipment.

[EHFR] h**ttp://en.wikipedia.org/wiki/Extremely_high_frequency** 30 to 300GHz. The short wavelengths permit better focusing, and are outside of most current communication bands, facilitating license assignment. Higher frequencies do not penetrate the atmosphere.

[HOTB] http://en.wikipedia.org/wiki/Hot_Bird_9 Hot Bird 9, a modern TV direct broadcast communication satellite.

[LOOP] http://launchloop.com The Launch Loop wiki website.

[MATW] http://matweb.com Lots of properties for commercial materials.

[RAZO] http://deepblue.lib.umich.edu/handle/2027.42/62283 Razor II variability tolerant computing design.

[SMAD] http://www.smad.com/about/smad3.html Space Mission Analysis and Design, Wentz et. al., 3rd edition.

[SMAL] http://www.researchchannel.org/prog/displayevent.aspx?rID=3427&fID=345 An excellent lecture video about the Terawatt Challenge by scientist Richard Smalley.

[SOLA] http://www.solarworld-usa.com/Made-in-USA-Our-produ.582.0.html Solarworld US plants

[SSKY] http://server-sky.com/RollControlV01 Roll from pitch and yaw:

[SSPS] http://gltrs.grc.nasa.gov/cgi-bin/GLTRS/browse.pl?2004/TM-2004-212743.html Space solar power satellites

[TRIQ] http://www.triquint.com/ GaAs radio chips

Open source software tools used:

http://www.libgd.org/ LibGD, a useful collection of 2D drawing primitives.

http://www.povray.org/ 3D modeling and ray tracing rendering engin

http://moinmo.in/ MoinMoin wiki software

http://www.math.union.edu/~dpvc/jsMath/ Rendering math on the web using javascript

http://wiki.themel.com/jsMathParser moinmoin plugin for mathgsl

http://www.gnuplot.info/ The gnuplot graph plotting package

http//www.gnu.org/software/gsl/ libgsl , The Gnu Scientific (math) Library

Math::GSL at any CPAN repository. Jonathan Leto's Perl implementation of libgsl

http://www.swftools.org/ SWFtools used for the website animations.

http://openoffice.org / it sucks, but less . It is open source, which helps.