SERVER SKY

Computation in Orbit

Keith Lofstrom **keithl@ kl-ic.com**

http://server-sky.com **2009 September 26**

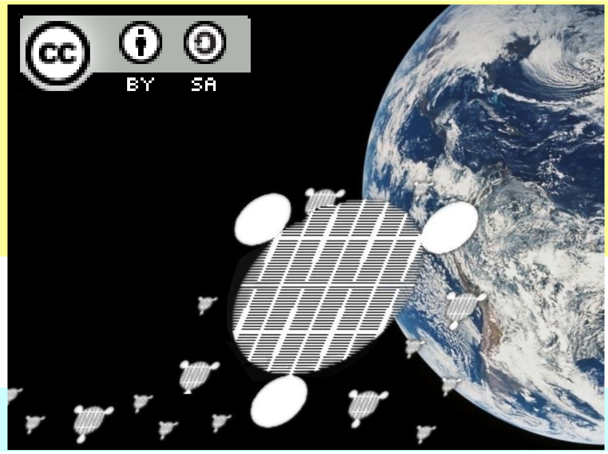

Abstract

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 black-body radiation. Arrays of thousands of server-sats form 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.

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 that total computing power for the planet doubles every year, implying that world computing energy demand doubles every two. We are not constructing enough clean power plants to meet this rapidly growing demand. Competition for the diminishing supply of power plant fuel will become increasingly deadly in the coming decades. The U.S. may have less generating capacity 20 years from now, while data center and data communication power usage increases to 40% of total load.

A 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.

Unless we learn from recent history, the actual outcome will be worse. During the California energy crisis, utilities reacted to high demand by shedding customers. Data centers are usually powered with

battery-backed uninterruptable power supplies, but these systems are limited, expensive, and inefficient. Data centers will shed compute load during blackouts, and go dark during long power outages.

Packets travel through dozens of switches between the data center and the end user. The internet is agile, and can route around failed links, but too many failed switches will result in inefficient routes, increasing the burden on the switches that remain. The result will be an increasingly slow, unreliable, and unpredictable internet. As "smart power" grids become 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 otherwise feeds the biosphere. Generating the world's energy needs (estimated at 40 Terawatts by 2050 [SMAL]) with solar cells will require millions of square miles of solar arrays. The 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 also requires huge amounts of infrastructure. Terrestrial solar energy is interesting, and useful for small and remote systems, but terrestrial solar is not a practical way to generate Terawatts of electricity.

The Sun fills space with unused energy. Space solar power satellites [SSPS] could capture some of this energy and beam it to earth. SSPS antennas produce intense microwave beams, focusing 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, requiring large rectennas.

Some SSPS power will drive data centers. However, the path from orbit to end usage is inefficient, with losses from transmission, sidelobes, power conversion, data center cooling, etc. A 20% efficient, one meter square solar cell in orbit intercepts 1300W of sunlight. Of the 260 watts of electricity it produces, perhaps 4% reaches the compute load in a data center. This does not include the power used to orbit the solar power satellite, repair it and supply thruster fuel, etc., lowering overall efficiency further.

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. A solar cell directly produces the high current. low voltage power that a modern 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 on the ground can be eliminated as well.

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 easier to make.

Server sky is many arrays of ultra-thin (100μm) 7 gram satellites. Each "server-sat" maneuvers by light pressure, and 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 dense stacks with conventional rockets to a 6411 kilometer altitude orbit. The server-sats 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 manufactured cheaply with existing factories, some idled by recent economic troubles. Although expensive to launch, they will be vastly cheaper than traditional satellites, and will quickly pay for their 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 cheap power can launch 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-quality 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. 10,000 server-sats can be stacked into a solid 1 meter column.

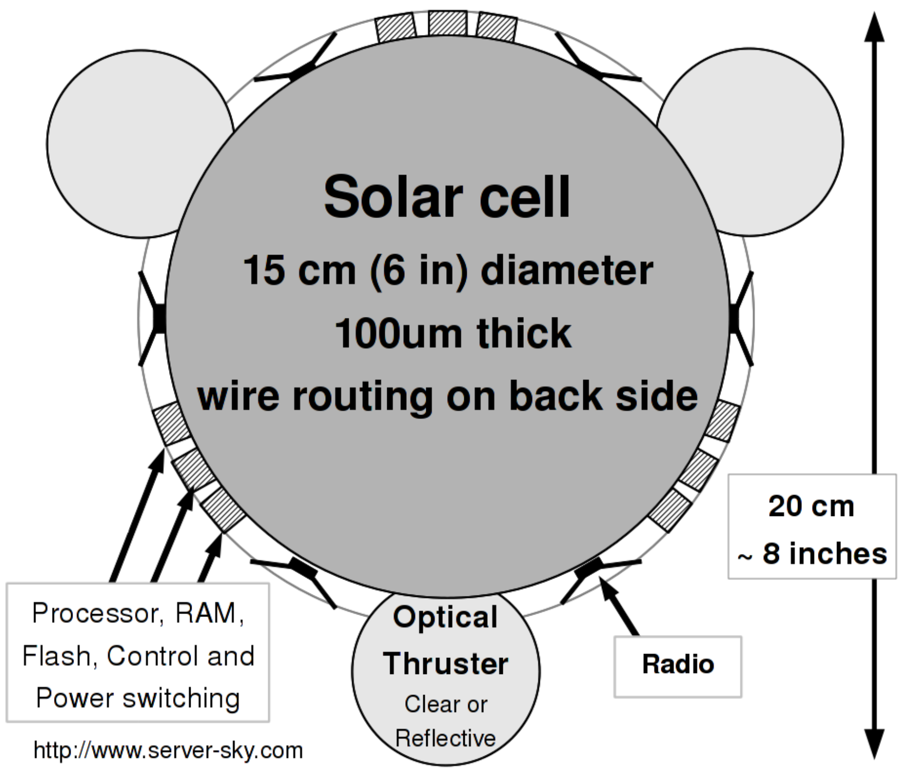

Decreasing server-sat weight reduces launch cost and results in a more effective solar sail. The current thickness target is 100 microns, though production silicon is often thinned to as little as 20 microns for

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 will become thinner than 5 or 10 microns, weighing less than a gram.

Everything on a server-sat is coplanar, for stacking during launch. The integrated circuit chips are arranged around the edge, and power is fed outwards (in separated zones) from the solar cell. If a portion of the solar cell shorts out or is otherwise damaged, the remaining circuitry will still work.

Electronics: Large databases will be distributed over many server-sats. A server-sat used for database or web service may need as much as a terabit of flash memory. 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 crystals for oscillators may be replaced by surface acoustic wave (SAW) devices. Processors and radios will be operated at low voltages and low impedances. 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 from strip-lines and other planar components.

Some components can be built into a very thin planar surface. Others cannot:

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

As server-sat manufacturing grows, we will see planar versions of “right side” components move to the left-hand column. There will be relentless pressure to make all components thinner. This will make server-sats lighter, easier to launch, and more maneuverable.

A server-sat uses a small array of radios (many more than the six shown) to communicate with neighbors

in the array, with other arrays, and with the ground. Server-sats 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 to 50N will be 70 milliseconds, less than U.S. transcontinental ping time through optical fiber.

Three optical thrusters perform pitch and yaw control for the server-sat. A server-sat normally points straight at the sun, with thrusters 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 thrusters are stimulated, center-line thrust increases. The thrust difference is small, but enough to keep an array of server-sats on station relative to each other. Roll control is indirectly, 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. 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.

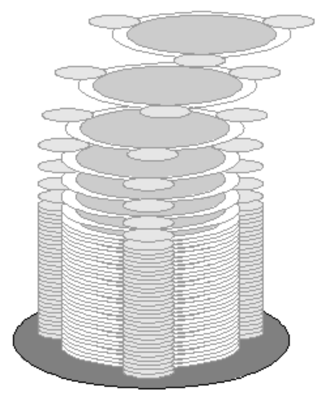

Server-sats are designed to have uniform thickness and matched mechanical properties, so they can be stacked in cylinders for launch. Some kind of very thin separator may be needed, or the server-sats may stick together with vacuum welding or vander Waals forces. If a server-sat plus separators is 110 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 lower m288 server-sat orbit permits a larger payload. 4200 kilograms is 600,000 server-sats and 3.6 megawatts of electricity. A server-sat array can produce more than 250 times the power (and radio bandwidth) of typical com-sats.

Light Pressure and Optical Thrusters

Server-sats maneuver by light pressure. Solar illumination is 1300 Watts per square meter at the Earth's distance from the sun. For absorbed light, the light pressure is the power divided by the speed of light, about 4E-6 N/m^2 or 4 microPascal. If the light is reflected, the pressure doubles to 8 microPascal. The pressure is tiny (sea level atmospheric pressure is 100 kiloPascals) but it is continuous. Light pressure pushing on a small, low-mass server-sat 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. 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 normally accelerate directly away from the Sun. 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 adds a small sideways thrust.

Server-sats get the most power if they face directly into the sun. They tilt to manuever. and that reduces the thrust. A 45 degrees sideways tilt produces 30% less power for computing and radio functions, which accommodate by slowing down. A 60 degree tilt produces half power (cooling significantly). Infrared light from the earth is mostly absorbed by the server-sat, and that creates some light pressure, too.

The version 2 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. An ideal thruster produces 16 nano-Newtons when fully reflective, and zero thrust when transparent. Real materials will show some reflection in transparent mode, and some transparency in reflective mode. Earth radiation (albedo and infrared) also reduces the effective thrust.

The thrust may vary between 4nN and 12nN (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 server-sat 10 degrees. Accelerating for 90 seconds, then decelerating for 90 seconds, turns the server-sat 60 degrees (not quite, as the thrusters are moving out of plane and become less effective when turned away from the sun).

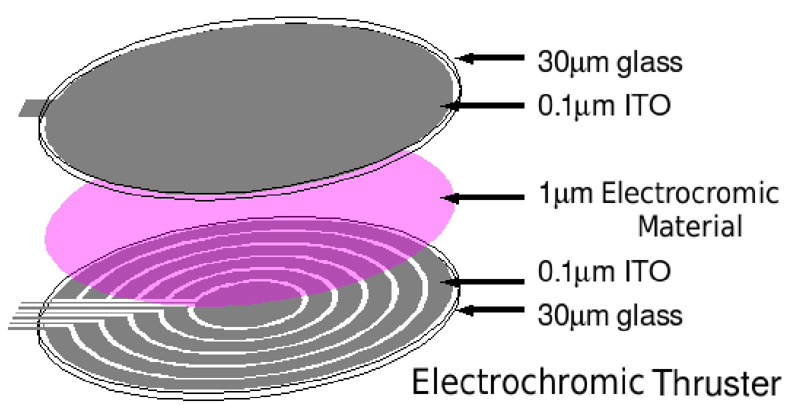

Each **Optical thruster** is made from two thin (30 micron) layers of commercial glass coated with of transparent Indium Tin Oxide (ITO) conductor on inner

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 as light and temperature controls for whole buildings.

Correcting for tidal forces: At the four "45 degree" points in the orbit, the server-sat is accelerated by tidal forces - the near end is pulled inwards by slightly more gravity and slightly less centrifugal acceleration, and the far end is pushed outward by slightly more acceleration and less gravity.

Optical thrusters are much more powerful than tidal forces, and 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 $0.29\,\mu rad/s^2$ for the m288 orbit, while the 5cm thrusters can provide angular accelerations of $70\,\mu rad/s^2$. This suggests a maximum mass-to-thruster ratio for server-sats: In the m288 orbit, the mass can grow to 1600 grams. First experimental server-sats can be thick and heavy. 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.

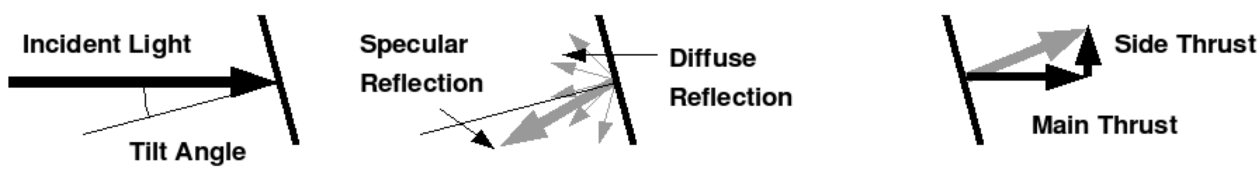

Thrust versus angle: Most of the area of a server-sat is a big solar cell, which absorbs most of the light that hits it. Some portion of the light reflects from the solar cell, and the reflections can be roughly divided into diffuse reflections (in all directions from the front side) and specular reflections (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.

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. Assume it is tumbling end over end , and the average resembles a sphere of the same radius (drag coefficient of 2). 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 1400km, about the altitude of Teledesic and Globalstar, 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

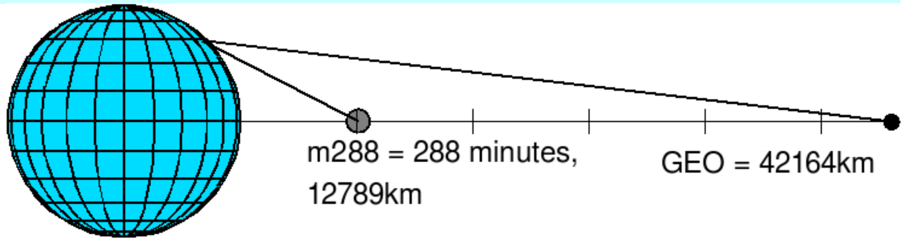

The first server-sat arrays will be deployed in a "4 hour" orbit, or more precisely a 23.9344696/6 = 3.989078 hour or 14360.7 second orbit (sidereal). This means it passes over the same spot on earth 5 times per day (= 6-1, the earth is turning underneath once per sidereal day), for a repeat time relative 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 the "gap" of the Van Allen belt, with an estimated unshielded radiation dose of 1Mrad/year [citation needed].

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. The orbital period should be an integer fraction of a day for logistic reasons.

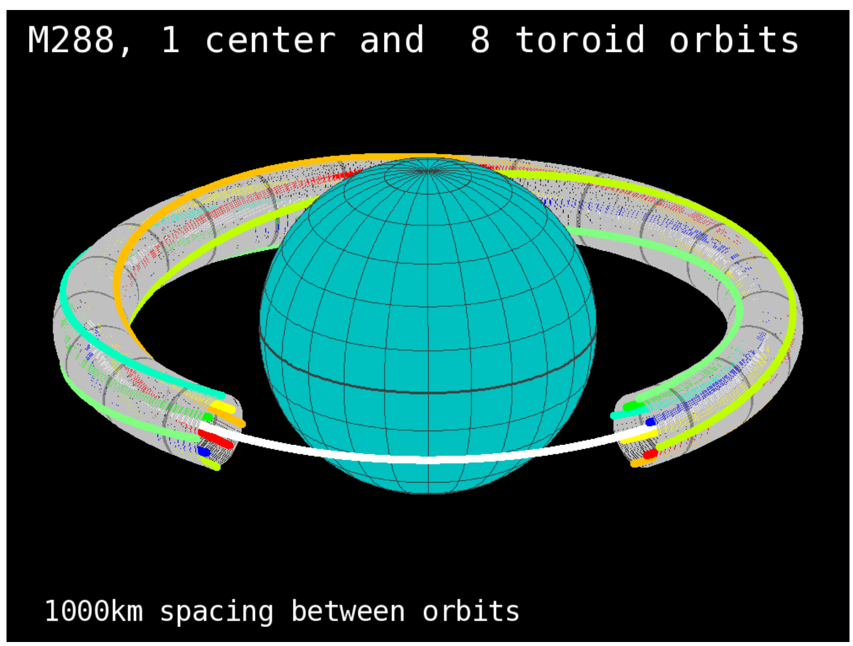

Defining Property Within the m288 toroid: Many orbits can be mapped on toroids surrounding a central orbit. If they all have the same semimajor axis as the m288 central orbit, they will have the same 288 minute synoptic period, and the same 14393 second sidereal period. If the orbits are mapped correctly, 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 m288 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 m288 central orbit can be seen at 58 degrees north and south latitude, at a distance of 10500 km. The round trip ping time to 58N 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. 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. This is still better than the 250+ millisecond ping time through a geosynchronous satellite, especially for cross orbit links to geosats below the horizon.

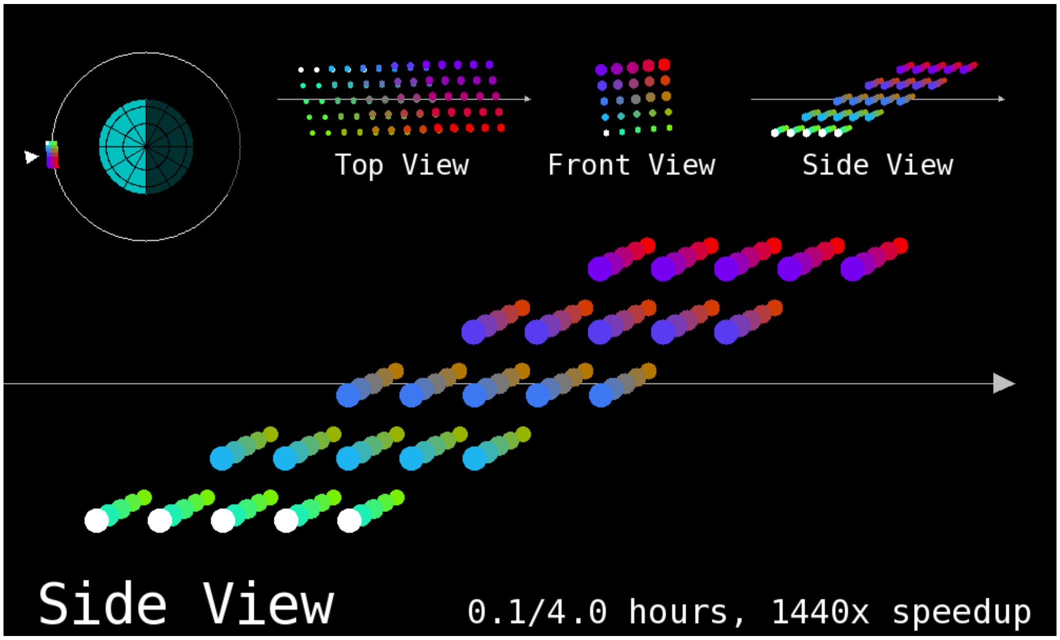

Arrays in toroidal orbits will distort, because different parts of the array are moving at slightly different speeds. Server-sats moving towards apogee are skewed towards it, by approximately twice the array dimensions. This is called apogee-skew and is shown here. See the server-sky website for an animation.

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. In the short term, there will be millions, not trillions, of server-sats. Still, they should be assigned orbits and positions compatible with a much more crowded sky. It is good to know, from the very beginning, that the region available for well behaved server-sky orbits can contain enormous growth.

Elevation above the horizon, interference with geosynchronous satellites: In the beginning, the first torus to be filled with arrays may have 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 tall 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 must relay through other satellites or landlines on the ground. Between latitudes 10 degrees north and south, the torus will be almost overhead, in the same band of sky as geosynchronous satellites. If 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 without interference.

Night Side Temperature: When a server-sat orbits into the night side of the earth it will only be heated by thermal radiation from the earth. At m288 in equilibrium, the night temperature of the server-sat will be around 127K (-146C). Higher orbits will have colder and longer “nights” and much longer “days”.

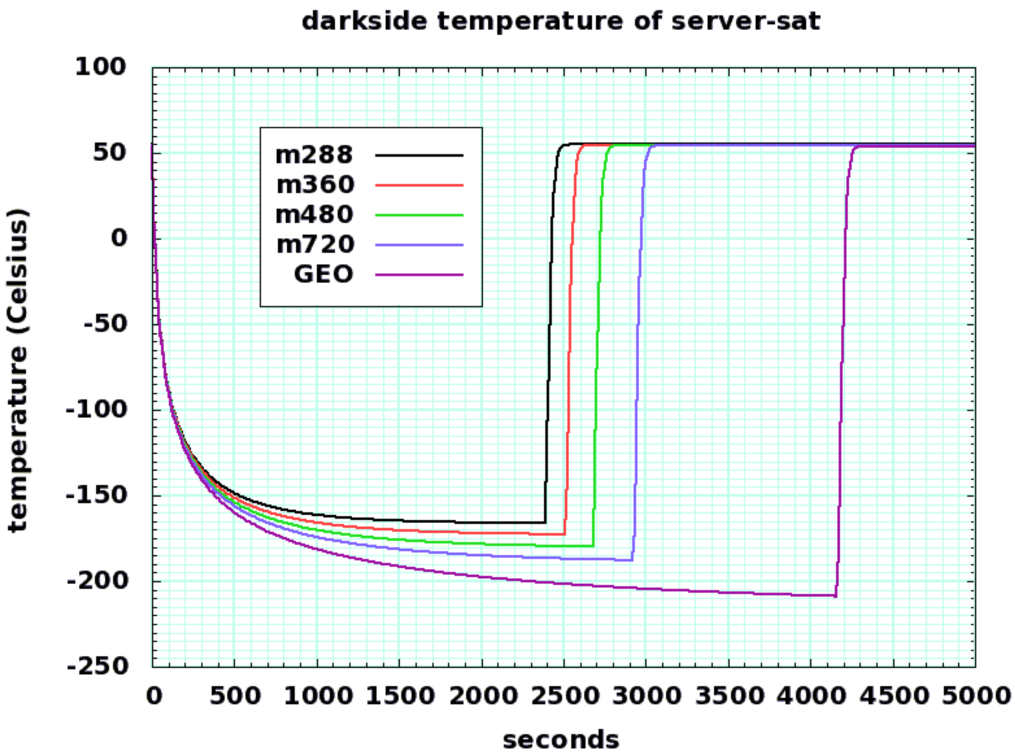

Shading

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.

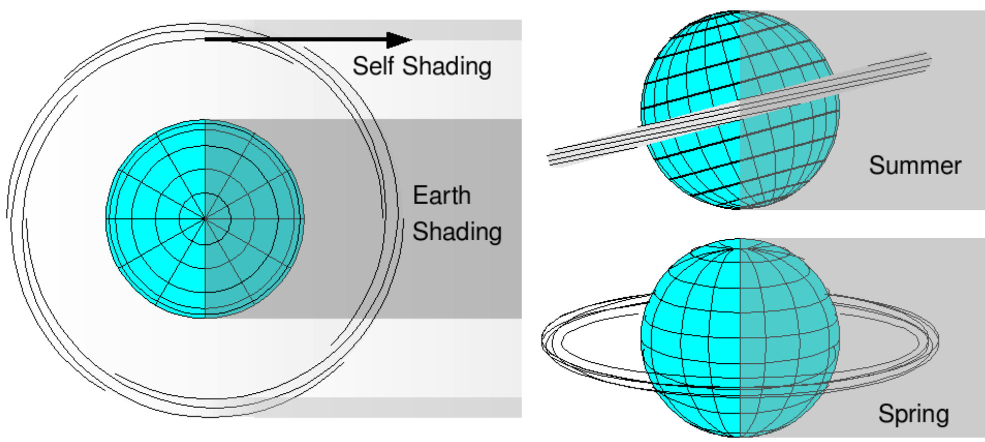

However, it is hard to imagine needing that many server-sats, even if they were operating mostly as space solar power satellites. Beaming about 5 watts each to the ground, this far exceeds the projected world demand for electricity. With a "mere" 10 trillion server-sats, the blockages would be 2% to 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.

Nighttime Illumination

Shading is not the biggest 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.

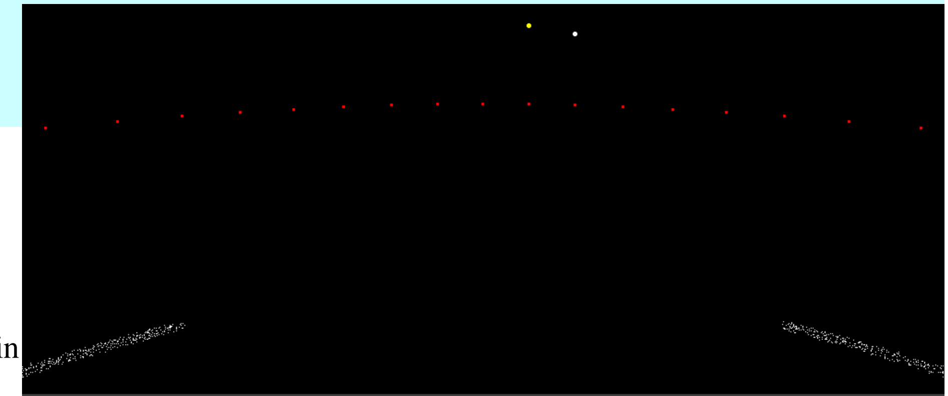

View of M288, geo, and moon

Server-sats reflect some light. Oriented directly at the sun, most of the reflected light goes 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. On the earth's surface 378000 km below, 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!

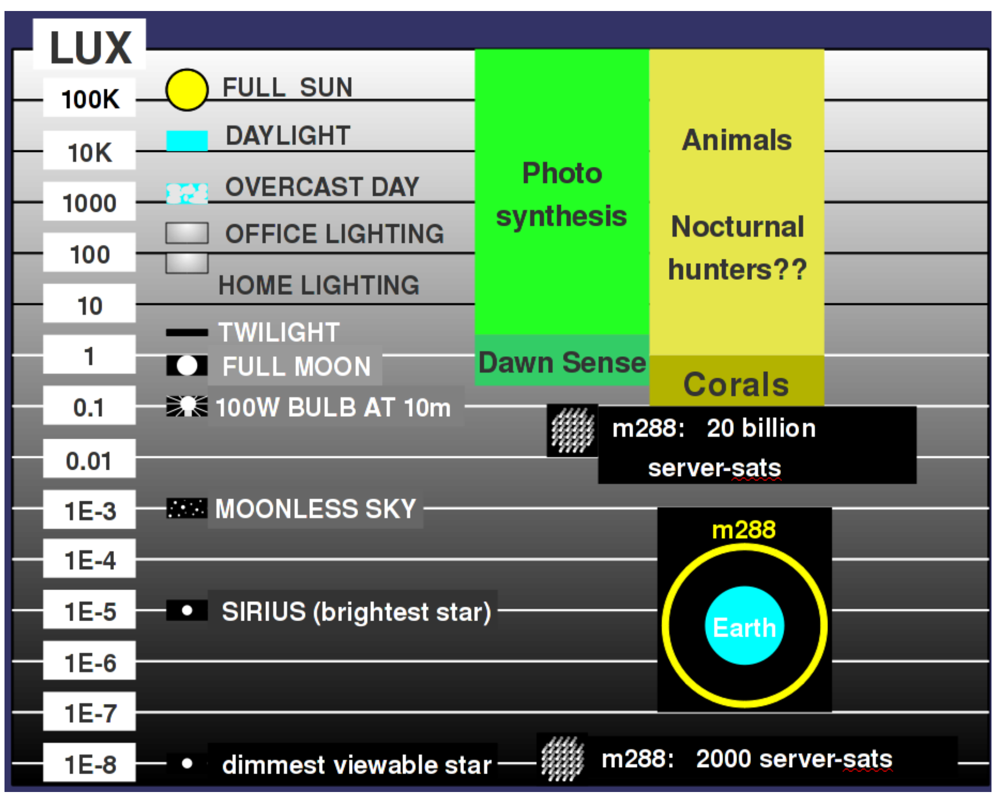

In normal circumstances, the array will be under control, and night-side server-sats will be oriented to reflect incidental light away from the earth. Front surface treatments will reduce night-time illumination further. However, circumstances change, civilizations collapse, and the server sats may stay in orbit for millions of years, long after we lose the ability to control them. In the worst case, micrometeroid bombardment may damage surface treatments, and they could tumble and scatter light in all directions.

This will be upsetting to astronomers, of course, but if civilization disappears, so do most astronomers. Protecting the biosphere is of supreme importance, however, and life evolved to adapt to existing cycles of day and night, full and new moon. Light pollution can alter this balance drastically, and perhaps destroy major ecosystems.

The total number of server-sats in at m288 will be limited by light pollution and its effects on **nocturnal animals.** For example, corals time their spawning to the monthly lunar illumination cycle, triggered by a fraction of full moon light. Interrupting that process with light pollution could kill the ocean reefs, vast swaths of ocean life, and might destroy the oxygen generation processes that keep us alive. Space computing and power is not worth the risk.

This may limit the near-earth constellation to less than 20 billion server-sats. This needs more study, with oceanographers, botanists, and astronomers helping engineers determine safe limits.

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, offset 60 degrees east or west. On average, each constellation of 5E12 server-sats will reflect 3E13 watts. At the same distance as the moon, each could shine with 2% of the brightness of the full moon, though reflectivity and attitude adjustments will greatly reduce this for properly controlled arrays.

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, opposite the moon, can provide full ground coverage, but will require constant station-keeping to stay in place.

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 of a completely predictable nano-gee space environment.

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. Assume a frequency of 38GHz and a wavelength of 8 millimeters. This is smaller than a server-sat, so each individual server-sat antenna array can direct radio energy into an angle of perhaps $\sin^{-1}(0.1)$ or 6 degrees, with a 600km ground spot.

Server sats in **phased arrays** have much better directionality. 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. For down-link, adding server-sats will improve spatial multiplexing bandwidth, with no practical limits on download bandwidth to billions of customers on earth.

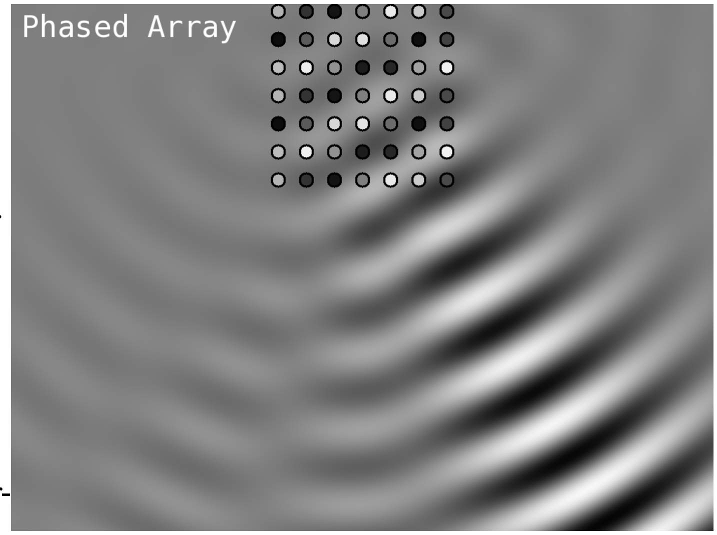

Phased arrays adjust the time delay of each server-sat transmitter so that the signals from each transmitter, located at a different distance from the receiver, all arrive at the receiver at same time. If each transmitter is emitting a pure sine wave, this can be accomplished by shifting the phase of 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 10,000 km 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 digitally 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 but slowly compared to round-trip ping time.

Array spacings are large compared to a radio wavelength, so they can scatter side-lobe energy into 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.

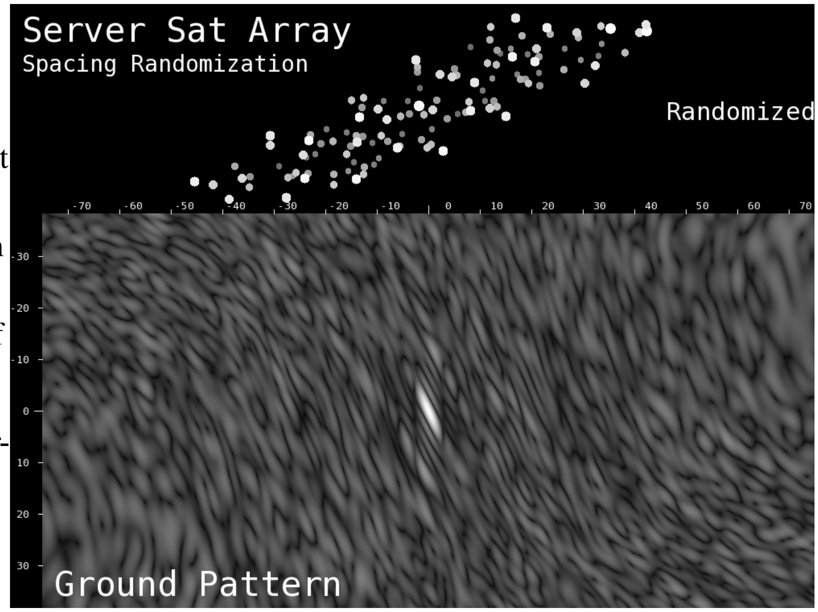

While the precision spacing of the server-sats permits accurate random spacing, that is probably not the best way to reduce grating lobes. It does show that we can do better than a regularly spaced lattice, and further study will find many sets of functions that will better reduce grating lobes.

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 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.

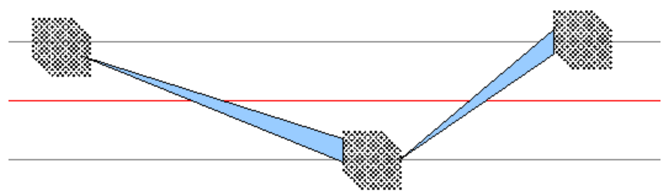

This can be partially ameliorated by using very narrow, **oblique**, and fractional-orbit beams between erver-sat arrays, such that the beams are pointed well above or below the orbital plane, and relatively weak when they reach GEO altitudes anyway. Space-to-space communications can use the 60GHz band, where the atmospheric absorption of oxygen is high. This will isolate the space links from potential jamming on the ground.

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. More delay will be added by the increased path length of the diagonal. If this is limited to 6%, then the beams (a fraction of a degree wide) will pass out of the ecliptic at a 19° angle . As the arrays get denser and larger, the beam size gets smaller, but the amount of inter-array traffic will increase faster than the number of communication paths.

Ground communication via other satellite services: Server-sky as described is confined to near the equatorial plane. This will restrict its usefulness for communications to ground sites with latitudes below 55 degrees or so, which precludes 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, server sat arrays can produce very tight beams with high power density. Off-angle reflections off small bits of space debris can be detected by many radar receivers optimized for the purpose. This permits 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 other typical orientation devices. 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 sats can estimate the sun angle from solar cell output. Surface gratings can be added to sense sun direction. An isolated server-sat will have some limited optical orientation capability.

A server-sat'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 measure 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 with the 60GHz intra-array communication links. If a server-sat measures phase within 1 degree at 60GHz, it can locate its many antennas with 20 micron accuracy. Improved software will improve measurement capabilities.

Server Sat Mechanical Behavior

Silicon is the construction material of choice - the solar cell is made of silicon, and the processors and memory are also. 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.

One possibility is to match mechanically to the **glass** rather than the silicon. 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^2s. Composite silicon-pyroC is 12.4 Kg/mm^2s, a much better match to glass than pure silicon, so standing waves at the interface are less likely. The composite has 50% better thermal conductivity than silicon, because pyro-C is an excellent thermal conductor.

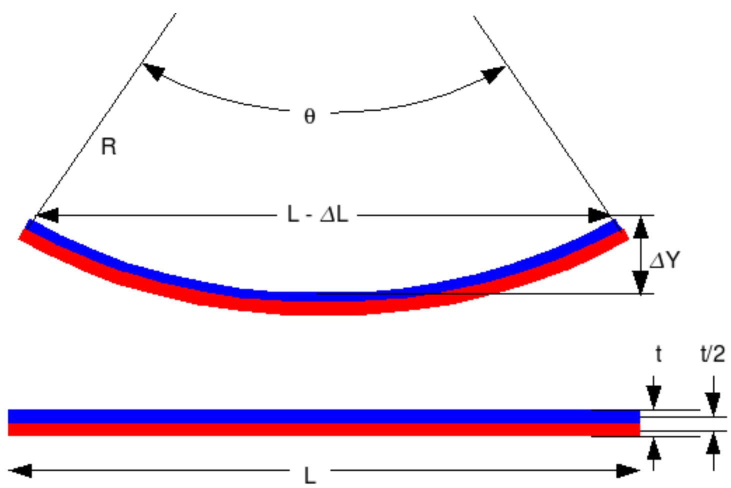

Curling can occur if the front and back sides of a server-sat (especially the solar cell) have different CTEs, and the server-sat undergoes repeated thermal cyclings. There is nothing inherent in a server-sat that establishes "flat" - it will flex until tensions and compressions are minimized.

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:

$$\Delta L \approx (1/6) \times (\beta/t)^2 \times L \qquad \Delta Y \approx (1/4) \times (\beta/t) \times L^2$$

Predicted curling as a function of size and other parameters

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

Radiation

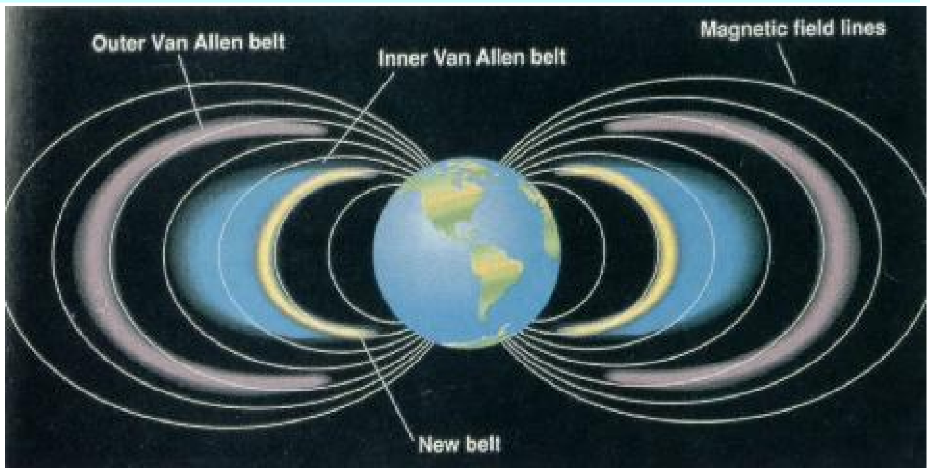

The M288 orbit (4 hour sidereal, 5 orbits per day relative to the earth) is located at a radius of 12789 km, about 6400 kilometers altitude. That places it between the inner and outer van Allen belts. This is a high radiation environment compared to low earth orbit. Ionizing radiation does nasty things to semiconductors: Latchup, Single Event Upsets (SEU, bit flipping), oxide charging , and flash memory errors. Traditional satellites are damaged by this much radiation, and there are few satellites operating in these orbits. Server sky resists these effects.

The space radiation environment is well characterized with computer models such as AF-GEOSPACE [AFGE]. Radiation effects are tested empirically in ground laboratories, at high dose rates, and extrapolated to years of space radiation exposure. Ionizing radiation does nasty things to semiconductors, including latchup, gate oxide charging, and Single Event Upsets (SEU, bit flipping).

Latchup: SCR paths 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. In some special circumstances, circuits called "voltage multipliers" are used to construct higher voltages for special needs. However, voltage multipliers are low current and easily switched off, so damaging high current latch-up will not be possible in these circuits, either. **Server-sats do not latch up!**

Gate Oxide Charging: 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.

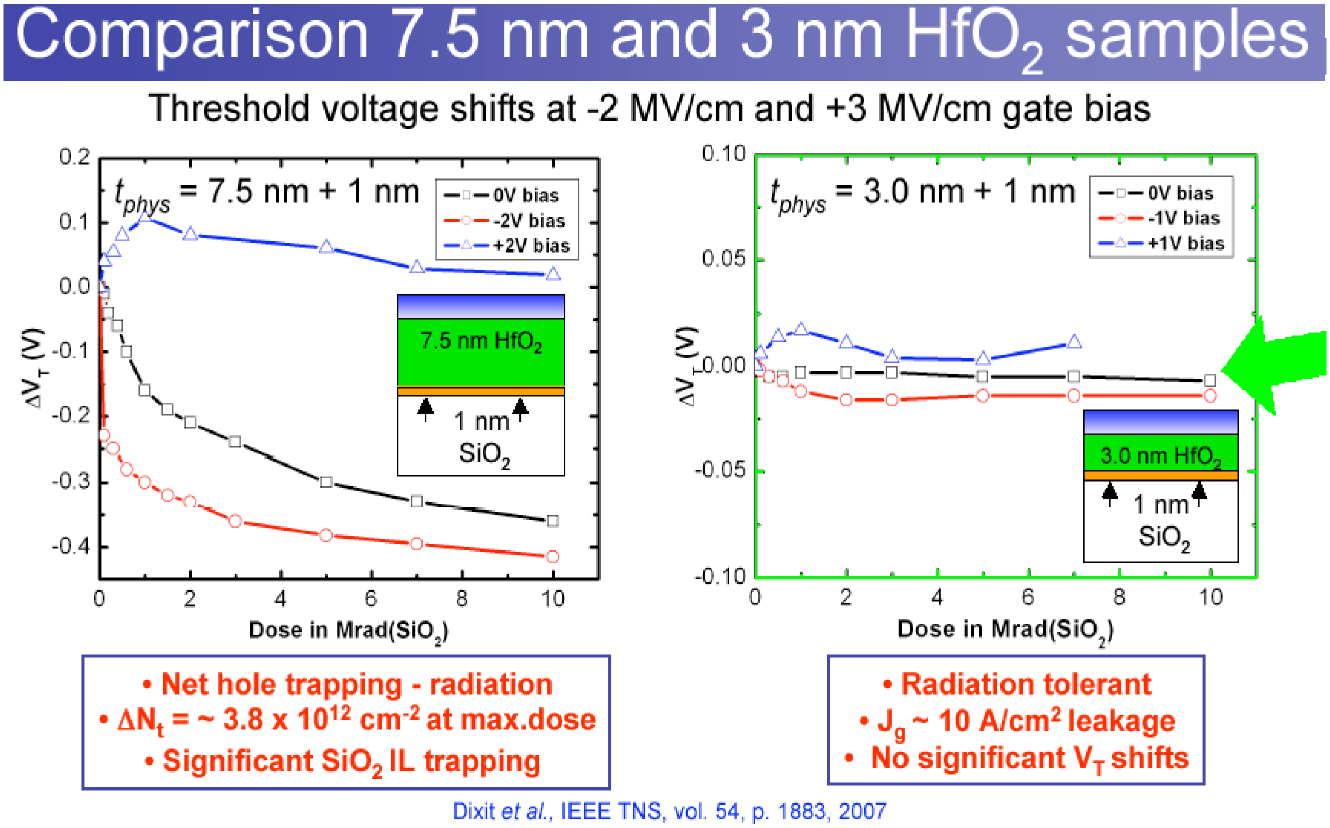

Single Event Upsets: The charge deposited by an ionizing particle can temporarily overwhelm a logic gate, or change the state of a register bit. This can cause incorrect computations. Many times, the cost is small. A calculation error in the I and Q signal of a software defined radio will just create a little extra noise. The cost is unacceptable when determining CPU state. Computation errors can be detected and corrected with redundancy, either by duplicating hardware or by repeating calculations, but this lowers performance and computational efficiency.

Razor error correction technology [RAZ] is 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. 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 errors occur. The performance improvement exceeds the cost of extra calculations, doubling overall performance. RAZOR technology will be common in microprocessors in a few years.

Radiation-initiated single event upsets may be detected and corrected by RAZOR-like technology. Since these SEU events may occur in the error correction logic itself, some redundant logic is needed.

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. 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.

Server-sats are cheap, and destroying one in a collision is unfortunate but not 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.

Server-sats can be maneuvered out of the way of accurately predicted collisions. Light pressure provides an unlimited supply of low thrust - in few hours, they can change orbit more than a kilometer away from a tracked impactor. Server sky will be deployed in orbits higher than most space junk. The vast majority of space debris is in lower orbits - it requires high launch velocities to even reach those altitudes.

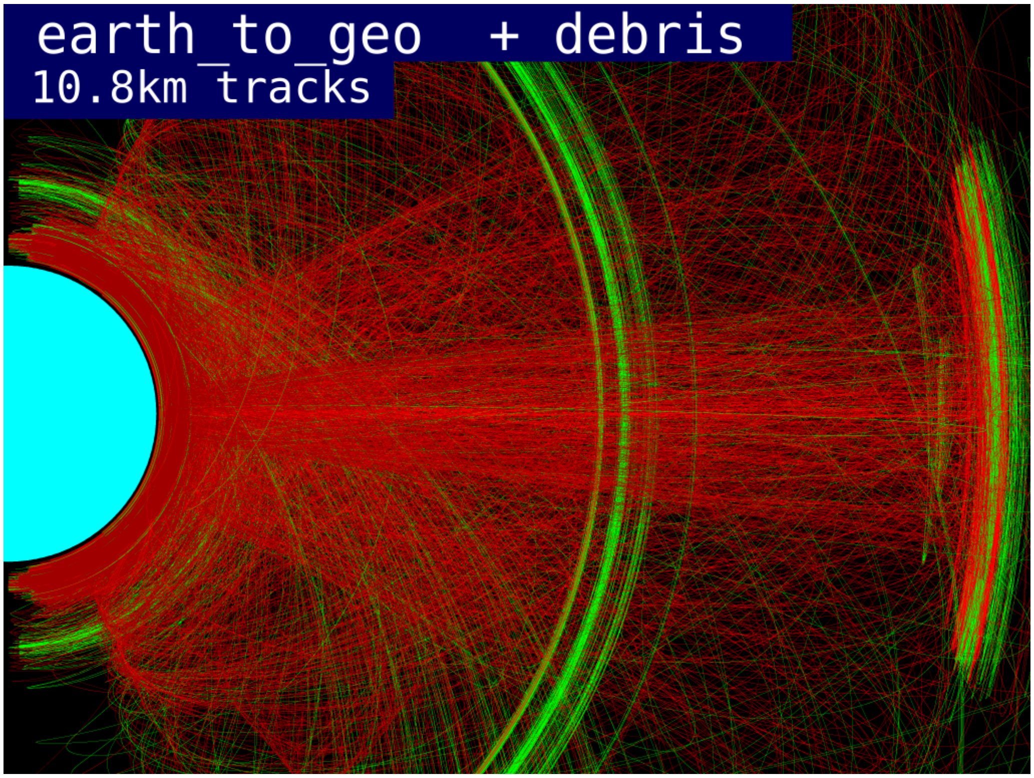

This plot shows objects tracked by NORAD in orbits up to geosynchronous altitude. This is an "HV" plot, with the horizontal H component representing the radius in the equatorial plane, and the vertical V component representing the distance above the plane. All orbits will look like an ellipse or banana 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 distances represented by the plot are 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 objects are small from centimeter scale (debris) to 100 meter scale (international space station). An "accurate" drawing looks blank.

Server sky version 0.2 will be located in m288, about one earth radius in altitude. The 100 kilometer white dot to the right of **this plot** 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.

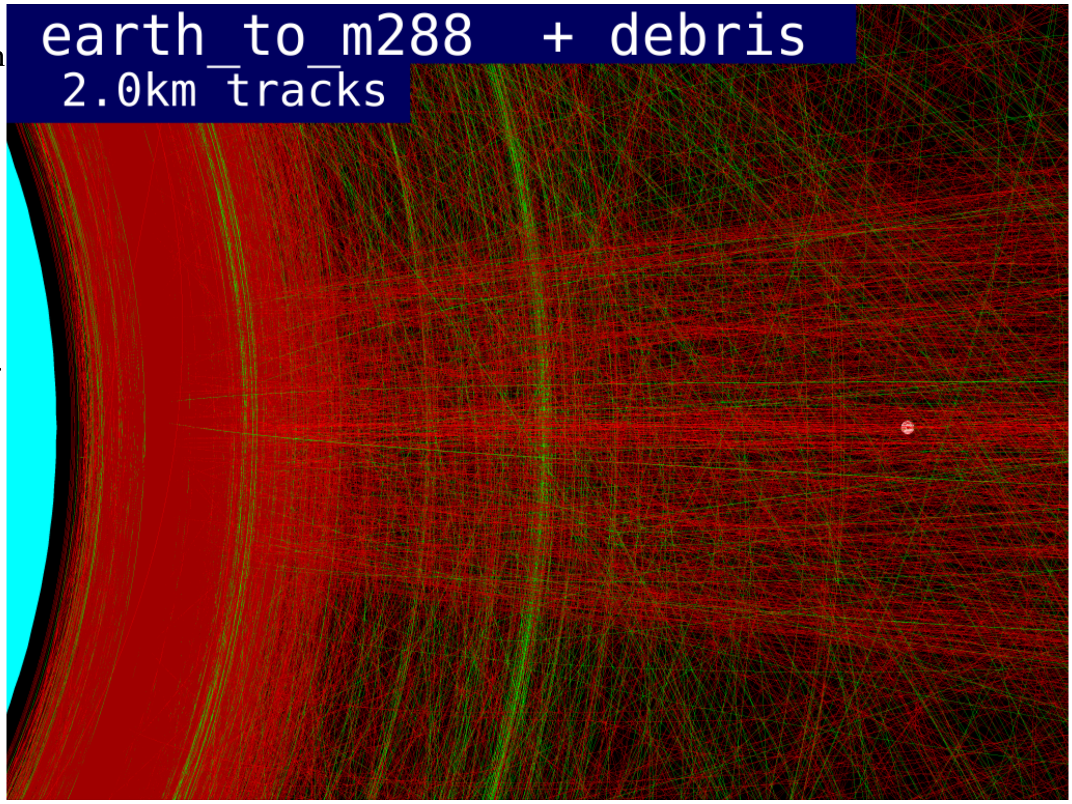

Except for the bands of navigation and communication satellites above, the density of drops off exponentially with altitude. The collision problem is not absent, but greatly reduced at altitudes above 1000km .

This plot shows just the m288 region. The square in the center of the plot is 10 kilometers on a side, as big as the city of San Francisco. If Server Sky orbits are threaded through the white spaces, they will *never* encounter one of the NORAD tracked space object..

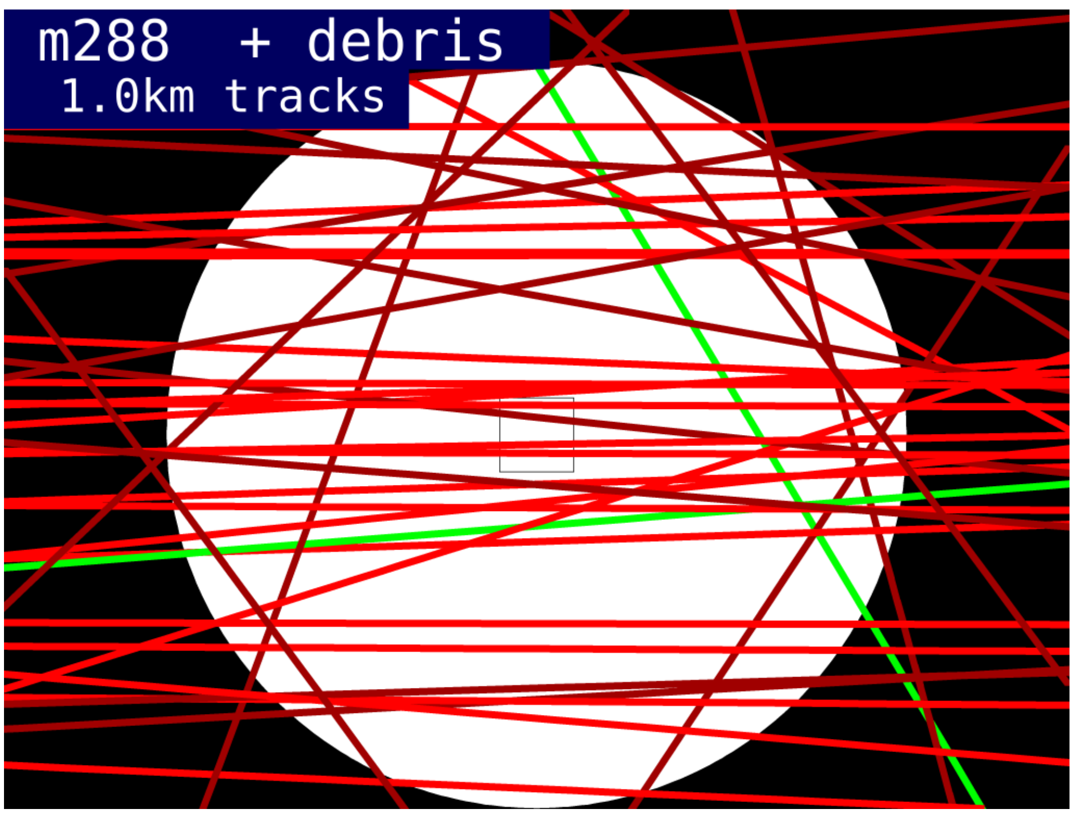

Only two potential assets cross the m288 server sky orbit: MOLNIYA 3-3 and BSAT 2B. The Molniya was launched in 1975 and may no longer be in use.

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.

NORAD does not track everything, especially small sub-centimeter objects difficult detect with ground-based radar. Orbiting radars can do better using very short wavelengths and large aperture antennas. Until that detailed mapping is available, centimeter objects that might destroy a solar cell or two on a large satellite may shatter a server-sat. Perhaps server-sats can be designed to survive a sub-centimeter collider passing through them. If a collider only hits one segment of the solar cell, the server-sat can keep functioning at lower power.

Possibilities

Long term speculation follows. We can bypass many limits to growth by moving more computation, communications, power production, and manufacturing 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.

That 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.
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 the principal constituents of lunar rock. It may be a long time before we can manufacture solar cells off the earth, and much longer before we manufacture deep sub-micron integrated circuits in space. But the materials on the moon are in a lower gravity well, and there are few ecological risks from using large amounts of lunar material to manufacture solar cells and glass substrates, 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, which is about as much energy as the finished product would produce in half an hour. Astoundingly high returns on investment are possible with lunar materials.

The same reasoning applies to asteroidal materials, applied to solar-system-scale server sky arrays.

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 space manufacturing. 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 megawatt-hour, may be possible someday.

Conclusions

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. The idea is in its early stages - you don't have to be a rocket scientist to contribute important new ideas, 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

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.

[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.

[BEKE] **http://www.thespacereview.com/article/131/1** Advanced Space System Concepts and Technologies: 2010-2030+

Ivan Bekey

[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.

[EHFR] **http://en.wikipedia.org/wiki/Extremely_high_frequency** 30 to 300GHz. Short wavelengths permit better focusing, and are above most communication bands, facilitating license assignment.

[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.

[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.

[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

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.