Online Journal of Space Communication

Volume 9
Issue 16 *Solar Power Satellites (Winter 2010)*

Article 16

Server Sky - Data Centers in Orbit

Keith Lofstrom

Server Sky - Data Centers in Orbit

Keith Lofstrom

Abstract

Server Sky is a proposal to build large arrays of 7 gram paper-thin solar-powered computer satellites launched into a 6400 kilometer altitude earth orbit. A server-sat is mostly a 100 micron thick, 6 inch solar cell, with processor memory and radio chips around the edges. Server-sats use light pressure for thrust and electrochromic light shutters for steering. Thousands of server-sats will be positioned in three-dimensional arrays, about 100 meters on a side. An array acts as a large phased array antenna transmitting thousands of communication beams simultaneously to ground receivers and to other arrays in space. In space, there is room for thousands of these arrays, with the total population limited by night sky light pollution.

A server-sat displaces 25 watts of ground-based electrical generation, cooling, and power conversion. A server-sat does not need the racks, cabling, power converters, land, buildings, and other infrastructure needed to build a ground-based server farm. These savings alone may pay for launch.

The space solar energy available for server-sats is practically unbounded. Server-sat arrays operate outside the biosphere, so the environmental impact of power generation and heat disposal is close to zero. In time, thinner server-sats, new launch techniques, and solar cells made from lunar materials may further reduce the environmental and economic costs of manufacturing and launch.

Thus, earth can return to what it is good at - green and growing things - while space can be filled with gray and computing things.

The Computing Energy Crisis

By 2011, traditional data centers will be consuming almost 3% of U.S. electrical power, and this fraction doubles every five years (EPA 2009). Computer technology improves, but the demand for computation increases more rapidly. Eventually, power efficiency improvements, and the power available for computation, will plateau. This limit on total computing will limit the availability of computers to solve critical world problems, such as the replacement of physical transportation by telepresence, or weather and ecological modeling.

As "smart power" grids become increasingly dependent on computing and internet communication to extract maximum efficiency from limited generation, the modern world may find itself experiencing deadly positive feedback loops, leading to cascading failure of the combined computing and generation grid.

A robust and green future may depend on computation that does not overload the electrical grid or the environment.

Opportunities

Space technologies lag modern consumer electronics, in some cases by decades. An advanced cell phone contains more transistors than most city-bus-sized satellites. Large satellites must survive the rigors of launch, but cell phones must survive the kilo-gee forces of hitting a floor. Satellites have tight power budgets, but they are lavish compared to the power available from a cell phone battery. Satellites are operated by large staffs of highly trained people, while some cell phones are operated by the oblivious. Satellites are produced by the tens and cost hundreds of millions of dollars to deploy, while cell phones are produced by the millions for less than one hundred dollars. The cell phone's advantages result from highly automated mass production based on photolithography, silicon VLSI technology, and the square-law increased value of large networks.

The bulk of a cell phone (and the bulk of a satellite) is mechanical packaging. The active area of a silicon chip is in the top five micrometers of the surface - the rest of the chip is a thick substrate, making the in-process silicon wafer mechanically robust in earth gravity. Modern chips are thinned after the wafer is complete, to improve wiring and heat removal. They are subsequently wrapped in many levels of packaging, mostly to shield the delicate silicon chips from moisture and mishandling. Thus, even a miniaturized cell phone is massive compared to the active silicon volume inside it.

The sun fills space with 384 trillion terawatts of unused energy. Space solar power satellites (Solar Power Satellites, 1981) will someday beam some of this power to earth, perhaps powering data centers through the grid. However, the energy path from orbit to CPU is inefficient, with losses from transmission, side lobes, power conversion and data center cooling. It is easier and more cost effective to use this power where it is generated. Indeed, if we want to someday exploit a large fraction of that 3.84E26 watts, we must do so on solar system scales, to both capture the energy and dissipate the waste heat.

The former director of advanced programs for NASA, Ivan Bekey, suggests 14 high leverage principles for space concepts (Bekey 2002). He suggests replacing structures with information, using coherent cooperative spacecraft to build distributed space systems, gossamer membranes, and transporting energy and information, rather than mass, through space.

Microgravity and highly predictable Newtonian mechanics permit large three-dimensional arrays of distributed satellites. Steady intense sunlight, high vacuum, and the 2.7K black body heat sink of deep space suggest ultra-light paper-thin satellites. Shielding, trusses, shaped dishes, and stored consumables are not practical at this scale. However, radiation, collision hazards, and the high cost of earth launch can be mitigated by the new possibilities. We need not construct spacecraft like we construct aircraft, and we can take advantage of semiconductor-style photolithographic manufacturing to print high volumes of low cost spacecraft.

Server Sky

Server Sky performs large-scale computation in space and radios data to and from earth. Solar cells directly produce the high current, low voltage power needed by modern CPUs. Direct radio communication with ground endpoints bypasses costly ground infrastructure. This increases the cost-effectiveness of space-generated power. Server sky arrays consist of thousands to millions of paper-thin 7 gram server satellites or "server-sats." A server-sat is a 15 cm solar cell with integrated circuits and optical thrusters attached on the edges.

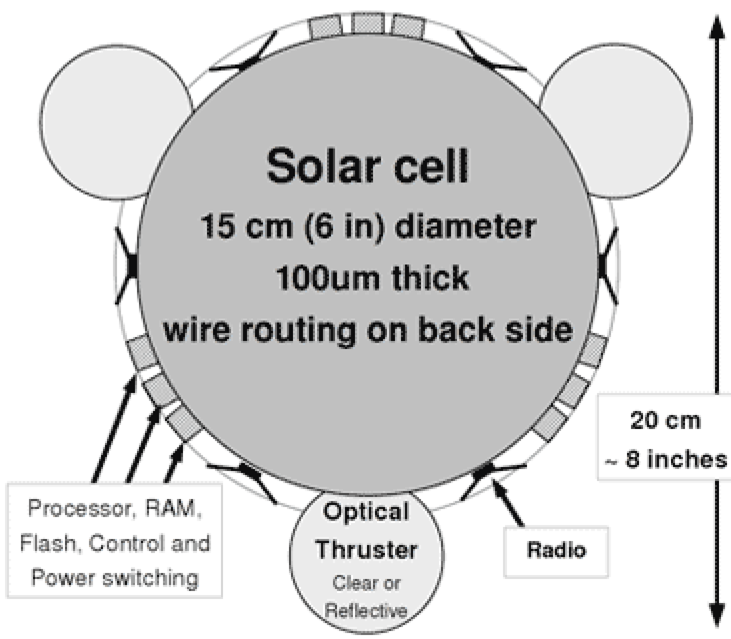

Solar cells, integrated circuits, and interconnect are two-dimensional. Modern IC die are thinned to increase thermal conductivity and reduce package height, and thin solar cells save material and cost. The first server-sats will be uniformly 100 microns thick. A estimated 10,000 server-sats can be stacked for launch in a solid one meter column, and deployed with a sheet-feed mechanism in orbit.

Maneuvering with Light Pressure

Server-sats maneuver by light pressure, about 4E-6 Newtons/meter2 for absorbed light. If the light is reflected, the pressure doubles to 8E-6 N/m2. The pressure is tiny (sea level atmospheric pressure is 1E+5 N/m2) but it is continuous. Light pressure pushing on a small, low-mass server-sat can add significant velocity over hours, days, and weeks.

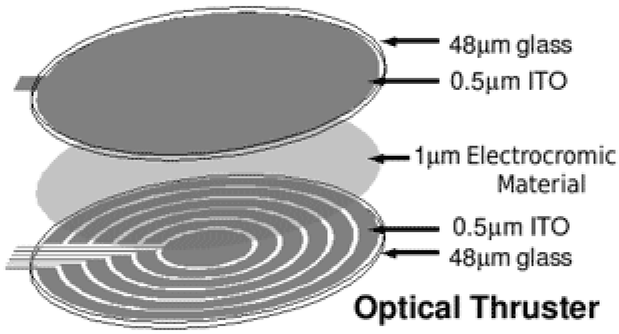

Three 5 cm diameter optical-shutter light-pressure "thrusters" at 120 degree angles around the periphery are electrically switchable from reflecting to transparent. Achievable thrust may vary between 4nanoNewtons and 12nN per thruster. The 24nN difference for three thrusters results in 3.4μm/s2 or 45m/hr2 of acceleration range, allowing large displacements over many orbits. A reflective thruster on one edge, while the other two are clear, produces a torque of 800 pico-Newton-meters, producing an angular acceleration of 14 degrees/min2. This can turn and stop server-sat 45 degrees in 3.6 minutes. Sideways-tilted server-sats can produce lateral thrust.

Deployment Orbits

Server-sat arrays will be deployed in constellations of 4-hour near-equatorial orbits with 14360.7 second sidereal periods. These orbits repeat relative to the ground every 288 minutes ("m288" orbits). The orbits will be chosen with slight inclinations and eccentricities to avoid collision and fill a toroidal region perhaps 100 kilometers across the minor dimension. The semimajor axis is 12789 kilometers, with an average altitude above the equator of 6411 kilometers. This is within the "gap" of the Van Allen belt, with an estimated unshielded radiation dose of 1Mrad/year. Equatorial orbits limit high latitude visibility, but minimize collision problems. Neighboring orbits will have small relative velocities. The chance of fratricidal collisions scattering material out of the region is small, even with millions of arrays containing thousands to millions of server-sats.

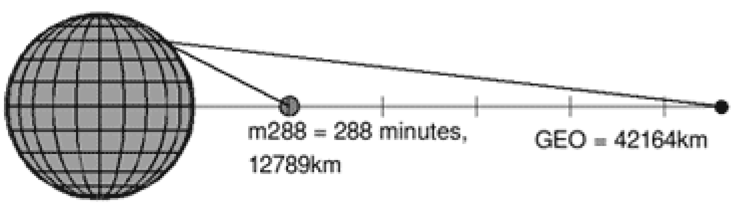

Server-sats shut down during the eclipsed portion of each orbit. At m288, the eclipse temperature drops to 127K (-146C). Higher, slower orbits have colder and longer "nights" and much longer "days."

Light Pollution

Server-sats have small but finite albedo and reflect some light. Server sky produces no visible daytime shading or glow. In daytime, even trillions of server-sats will have no noticeable shadowing effect on earth or on each other.

At night, the sum of the diffuse reflected light from billions of server-sats may appear as cloudy light in a thin band (about as wide as the full moon) around the equator, interrupted by the earth's shadow. Normally the array will be under control, orienting night-side server-sats slightly off-axis to reflect light away from the earth. But there will be some reflected light, astronomically obscuring dim objects in the direct path of the array. If the arrays are "bunched" into portions of the orbit, then there will be large dark gaps in between brighter spots, and astronomical imagers can be shuttered while the bright spots are in view.

Surprisingly, the effect of light pollution on biological systems sets the most important limit on server sky density (as well as a limit on any other kind of large area system in geosynchronous orbit or below). Unmanaged and tumbling server-sats may stay in orbit for thousands of years after a hypothetical collapse of civilization, scattering light in all directions. Life is adapted to existing cycles of day and night, full and new moon, and light pollution can alter this balance drastically. For example, corals spawn synchronously to the monthly moonlight cycle. Confusing that process with light pollution could kill coral reefs and vast swaths of ocean life. The maximum permissible light pollution from a worst-case control failure, assumed to be around 3% of full moon illumination, and 30 times a moonless midnight sky, limits the m288 constellation to less than 100 billion server-sats. Larger populations should be in higher orbits.

Radio

Multi-band server-sat radio chips have many low-power inputs and outputs, and connect to many planar antennas. Radios provide femtosecond-precision average array timing, micron-precision server-sat location and orientation within the array, and orientation to other arrays and to GPS and ground systems. Accuracy derives from continuous monitoring and averaging, differential and quadrature analog signal processing, and the ultra-low vibration and perturbation of a highly predictable nano-gee space environment. Thermal drift may be high, but temperatures can be accurately measured for tuning compensation.

Server Sky is blind; it has no star trackers or gyros. Server-sats can estimate the sun angle from solar cell output. Surface gratings may sense sun direction.

A server-sat's main "sense" is radio. It shares timing information with neighboring server-sats, permitting precision orientation and location computations. 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 does the same with 60GHz intra-array communication links. If a server-sat measures phase within 1 degree at 60GHz, it can locate its many antennas relative to the array with 20 micron accuracy. Continually upgraded software improves measurement capabilities.

As a new service, Server Sky may be allocated EHF frequencies for the down-link. Assuming a frequency of 38GHz and a wavelength of 8mm, each individual server-sat antenna array can direct radio energy into an angle of about 6 degrees, for a 600km ground spot. Server-sats in sparse three-dimensional phased arrays produce tighter beams than single devices. Arrayed transmitters can beam the sum of many different phased signals to sub-kilometer ground spots.

Server Sky operates near the equatorial plane, precluding communications to ground sites above 55 degrees latitude, including much of northern Europe. Existing services like Iridium, Globalstar, TDRSS, and many satellites in GEO may be used as relays. As traffic grows, this may provide extra revenue for these services, while not competing with their latency-sensitive telephony markets.

Developing countries near the equator may find server sky a low cost alternative for data centers and internet service, obviating the need for expensive fiber optic deployment and high reliability power systems.

A Server Sky array may be used as a narrow-beam radar transmitter to locate space debris and other satellites. Working in conjunction with radar receiver satellites and ground stations, server-sat arrays produce tight beams with high power density.

Server-Sat Mechanical Behavior

The solar cell is silicon, as are the processors and memory. The server-sat undergoes wide temperature changes when it passes in and out of shadow. Reliability improves when the non-silicon portions are made of composite materials that match silicon's 2.6E-6/K coefficient of thermal expansion (CTE).

Server-sats need transparent materials and conductors that closely match silicon mechanically. Metals have high CTEs, while SiO_2 has a low CTE. Slotted metal wires with SiO_2 in the gaps make a "material" that is both conductive and has the same CTE as silicon. This is common for stress reduction on integrated circuits.

Curling occurs if the front and back sides of a server-sat have different CTEs. Server-sats may not stay flat, but will flex until tensions and compressions are balanced. Curling can be measured and characterized and radio phases adjusted to work with curled surfaces. Careful choice of materials and layer thicknesses minimizes curling.

Radiation

The m288 orbit is located between the inner and outer van Allen belts, where ionizing radiation damages semiconductors and data. The effects include latch-up, single event upsets (SEU, bit flipping), oxide charging and flash memory errors. The high radiation damages traditional satellites, and few satellites operate in these orbits. Server sky resists radiation.

CMOS circuits latch up by the activation of PNPN parasitic paths. These paths cannot activate if the supply voltage is less than a diode turn-on voltage. Solar cells produce less than a diode voltage, so server-sats do not latch up!

Silicon dioxide develops a positive charge when irradiated, causing damaging threshold shifts in CMOS transistors. Hafnium oxide traps electrons and develops a negative charge. Recent work by Dixit (2008) shows that a combination of both materials produces a radiation-hard gate oxide stack, withstanding 10 Megarad from a Cobalt 60 source with minimal threshold shifts.

Charge deposited by an ionizing particle can cause SEU, temporarily overwhelming logic gates, or changing states in memory. RAZOR technology (Shidhartha 2009) reduces noise margins while adding error detection circuitry, repeating calculations when errors occur, which doubles overall performance. RAZOR technology will soon be common in commercial microprocessors. While the timing information driving RAZOR is not the same as radiation-induced node charging, we may be able to detect ionizing particles by other means, such as scintillator layers above and below the active circuitry, computational checksums, grey code counters, etc. Triple redundancy can be used for some circuitry, but complete triple redundancy wastes too much power.

Server-sats have low thermal mass and low lateral thermal conductivity. Power may be diverted to heating wires behind small regions of the solar cell, or to individual chips, heating them to 300C or higher. The temporary high temperatures anneal trapped charge and lattice displacement damage. This is difficult to do with traditional satellites.

Space Debris Collision Hazards

Server Sky must not be permitted to come close to other expensive space assets. Misbehaving server-sats must be captured, compacted, and de-orbited before they can damage other satellites. Fortunately, there are no valuable permanent space assets orbiting near m288, just debris.

Centimeter scale or smaller impactors will punch holes in server-sats. With proper redundancy and crack-stopping design, small impactors reduce capacity but won't disable server-sats. Partly crippled server-sats can still maneuver out of their array to collection areas for disposal.

Server-sats are cheap. Destroying one in a collision is unfortunate but not expensive. The main cost of a server-sat collision is more space debris, which can damage other server-sats. Traditional large satellites encountering server-sat fragments may lose a solar cell or two on the skin, but won't be destroyed. Server-sat debris will resemble natural micrometeorites rather than large satellite debris.

Server-sats can maneuver out of the way of accurately predicted collisions. Light pressure provides an unlimited supply of low thrust - in a few hours, they can maneuver kilometers away from tracked impactors.

Possibilities

Big compute jobs like weather prediction or animation rendering are not real-time and can tolerate high communication latency. Constellations of trillions of server-sats in the Earth-Moon Lagrange positions are 60 times further away from earth than m288, so they produce 1/3600 of the worst-case light pollution.

Later generations of server-sats may become "compute-light" and "transmit-heavy," beaming the power as microwaves to rectenna arrays on the ground, making power for the electrical grid. These will be a much lighter and more easily deployed version of a Space Solar Power Satellite (SSPS), potentially delivering 10 kW per kilogram of earth launch. These should be deployed at Lagrange positions, again to minimize light pollution.

A server-sat is mostly silicon, glass, and aluminum, similar to lunar rock. Advanced technology integrated circuits require mega-factories on earth, but low-tech lunar-manufactured solar cells are a possibility. A cubic meter of lunar regolith might provide materials for 5 million 300 milligram solar cells, which could be mated to earth-manufactured integrated circuits in small automated assembly factories in orbit.

Septillions of server-sats made from asteroidal materials can generate trillions of terawatts of usable electric power for computation and space manufacturing, millions of times the solar power received by the earth. Server-sats near the asteroid belt at 2.6 A.U., receiving 200 watts of sunlight per square meter, have an equilibrium temperature of 200K. This can heat the earth. If we limit the heating to 0.5C, the average deep sky temperature can rise from 2.7K to 84K. We can capture 2% of the sun's output, 8 trillion terawatts.

Server-sats between Saturn and Uranus, at 15 A.U., have an equilibrium temperature of 84K. We could fill the sky with them, capturing all of the sun's output. The lower temperature also permits lower energy per bit computation and communication. The technologies for such advanced server satellites, and the computation jobs that justify them, are left to the imagination of the reader.

Conclusion

Server Sky may be the near-term commercial application that pays for the permanent expansion of life and civilization into space. Server Sky can remove large scale computing and power generation from the biosphere, leaving the earth safe for green and growing things, while filling space with gray and growing things.

REFERENCES

1. Bekey, I. (2002). Advanced Space System Concepts and Technologies: 2010-2030+, The Aerospace Press, El Segundo, CA.
2. EPA Final Report to Congress - Executive Summary. (2007, August). Retrieved from http://www.energystar.gov/ia/partners/prod_development/ downloads/EPA_Report_Exec_Summary_Final.pdf.
3. Dixit, S., Feng, T., Zhou, X., Schrimpf, R., Fleetwood D., Pantelides, S., ... Feldman, L. (2008, May). Interface Structure and Charge Trapping in HfO2-based MOSFETS. MURI Annual Review. Retrieved from http://www.isde.vanderbilt.edu/content/muri_2008/dixit_muri 2008.pdf.
4. Shidhartha, D. (2009). Razor: A Variability-Tolerant Design Methodology for Low-Power and Robust Computing. PhD Thesis, University of Michigan. Retrieved from http://deepblue.lib.umich.edu/bitstream/2027.42/62283/1/ siddas_1.pdf.
5. Solar Power Satellites (1981). Office of Technology Assessment, Washington D.C.
6. Lofstrom, K., et. al. Server Sky publicly editable Wiki website. http://server-sky.com/.