Tesla, SETI & Globus: Understanding the Digital Ribbon.

Part I

"Ever since Benjamin Franklin flew a kite with a key attached, the dream of tapping such power has lit the darkest corners of mans imagination."

Thomas Edison vs. Nicola Tesla: The battle of AC & DC power

In 1890 a battle was waged between the well-known Thomas Edison and the much less renowned Nicola Tesla. Their dispute was over Direct Current (DC) and Alternating Current (AC), and its outcome would change the world forever.

On Sep 4 1882 Thomas Alva Edison opened the first ever DC power station on Pearl Street in New York City. Its consumption catered to Wall Street, some wealthy private residents, and to shop owners as a novelty.

Around the same time, budding young genius Nicola Tesla experienced an epiphany on "how to make AC current work." With his new insight, Tesla sought out Edison and attempted to convince him of the potential of AC.

Edison had a large vested interest - both financial and emotional - in the DC power plants that he had been building, and which the "robber baron" J. Pierpont Morgan had been financing. Unfortunately, because DC power was only available to consumers located within a mile of a DC power station, it was limited in its ability to power areas that could not support its costly equipment. In addition DC power sources suffered from frequent hardware breakdowns and limited ability for improvement.

Tesla, on the other hand, backed the concept of AC current. Tesla's designs introduced the concept of placing generating stations far from consumption centers. The systems were much more reliable and scalable. AC power allowed energy to be produced in one location and sent as far as twenty miles before utilization. The only drawback was the potential for electrocution.

Tesla eventually found funding from a man named Westinghouse, who owned patents in DC current. Westinghouse was convinced, despite his investment in DC, that AC was a superior, cost-effective system and thus a wiser investment. He became Tesla's major backer and purchased his patents.

A fierce battle ensued; Edison went so far as to have public demonstrations of the dangers of AC current on stray animals. J.P. Morgan battled Westinghouse in a different manner, spreading rumors among Wall Street investors that Westinghouse's finances were unstable. Investors began to shy away from providing Westinghouse with the capital he needed to implement AC power.

In the end Tesla's superior technology won. The cost was great; Westinghouse was almost bankrupted and Tesla gave away concessions to all future earnings on the patents.

What if Edison triumphed? Could you imagine the world today if every square mile of development requires its own power plant? How many people would be with out such a basic utility? What devices and innovations would never come about? How high would the cost of electricity be?

It's begs the question, "Should only the wealthy have such resources?"

Part II

SETI, seeds & the upcoming revolution

"There is a digital ribbon of ones and zeros circling the globe and intertwining our lives."

SETI foundations of the future

The Search for Extraterrestrials Intelligence (SETI) was faced with a problem of too much information. The project had huge amounts of data from state of the art telescopes but lacked the budget needed for a supercomputer large enough to break down the information.

"In 1994 David Gedye and Craig Kasnoff, two computer scientists from Seattle, came up with a brilliant solution to this problem: instead of running a single large computer for a long time, why not run thousands of small computers for short periods of time?"

This could be done, they reasoned by having people from around the world download a simple data-processing program to their personal computers. The program would not interfere with the regular operation of the computer, and would only run during its idle time. Once a batch of data was processed, the computer would return it and receive a new batch to process in its place."

SETI's solution took into account the growing processing power of personal computers (PC) and the fact that processing power of those PCs is unused the vast majority of the time. Using these facts and an innovative program that took advantage of processor cycles while a PC's screen saver was active, SETI changed the face of computing forever.[1]

The methods used to implement SETI's goals, which became known as SETI@home, are known as the distributed computing model, the first step towards being able to use computing power as a pure utility similar to electricity. With many millions of home and business PCs comes a vast wealth of resources ready to be tapped into in a manner similar to how dams tap rivers and windmills tap the wind. Unlike rivers or the wind, this body of resources will only grow as more people and businesses acquired faster computers. An individual or organization that need to perform large amounts of data crunching will no longer be forced to buy large expensive computers; they can simply utilize a plethora of distributed resources to solve their problems. Such is the birth of the *utility model* of accessing computational resources.

Seeds – Half the problem.

Great! We can tap this new computing resource to solve every computational problem, right? Wrong! Anyone reading this article that has a background in software or research knows this approach will not work for all problems.

"Some problems must be divided into myriad minuscule tasks; because these fine-grained problems require frequent internodes communication, they are not well suited for parallel processing. Coarse-grained problems, in contrast, can be divided into relatively large chunks."[2]

There are basically two primary types of problems. Envision seeds, some coarse and some fine. Software, like seeds, can be broken down into pieces. Coarse tasks work great for distributed computing in that they can be split up into many pieces and shipped out to millions of small computers. Fine tasks, on the other hand, require frequent intercommunication between nodes (or computers). These fine tasks need to frequently communicate their piece of the problem to other pieces of the problem to create a solution. The time between communications introduces the problem of latency. To solve that problem we introduce the second piece, one that is shaking the foundations supercomputing. "Beowulf"

---

[1] As of April 2002 more than four million users have downloaded the program onto their personal computers! This makes SETI@home easily the world's largest supercomputer.

[2] - The Do it yourself Supercomputer. Scientific American: Aug 2001, William W. Hargrove, Forrest M. Hoffman and Thomas Sterling.

Beowulf, a cluster of Linux PCs on a high performance network, was born in 1994 at the NASA Goddard Space Flight Center. The triple punch of advancing PCs, innovations in high performance networking, and the Linux operating system has created a new generation of supercomputers that rival the cost to performance ratio of anything on today's market.

"In the middle of producing Star Wars Episode II, ILM switched from using Silicon Graphics RISC-Unix workstations to Intel-based Dell systems running Linux. The company deployed 600 Pentium 4 workstations that were about three times faster than the machines they replaced yet a fifth of the price."

-Cliff Plumer, chief technology officer at special effects house Industrial Light and Magic.

Marrying the technologies of distributed processing and next generation cluster PCs creates the underpinnings of the complete utility model for computational resources. Tapping into this digital ribbon will alter the future of supercomputing and radically change the way research and innovation is pursued.

These disruptive technologies are so far reaching in their influence that a recent report by the National Science Foundation (NSF) declared their presence to be "a once-in-a generation opportunity to lead the revolution in science and engineering"

Advanced Cyberinfrastructure Program - The upcoming revolution

The final report from the Blue Ribbon Advisory Panel on Revolutionizing Science and Engineering through Cyberinfrastructure stated:

"The panel's overarching recommendation is that the NSF should establish and lead a large-scale, interagency, and international coordinated Advanced Cyberinfrastructure Program (ACP) to create, deploy, and apply cyberinfrastructure in a way that radically empowers all scientific and engineering research and allied education. We estimate that sustained new NSF funding of $1 billion per year is needed to achieve critical mass and to leverage the coordinated co-investments from the other federal agencies, universities, industry, and international sources necessary to empower a revolution. The cost of not acting quickly or at a sub-critical levels could be high; both in opportunities lost and increased fragmentation and balkanization of the research communities."

The report goes on to state how critical these emerging technologies are to research and engineering. It discusses how these technologies placed together with software tools will create a cyberinfrastructure that different research communities as a whole can leverage as a single resource that is 100 to 1000 times faster and more effective than are available locally.

The ACP, like Tesla, will shift the dynamics by attacking several key points.

1. The practice of placing generating stations far from consumption centers. This allows for more flexibility in consumption and opens the door to both economics of scale and giving the power back to the community in the most efficient manner possible.
2. Modern supercomputers are the equivalent of DC power, technology that has forced itself into a corner and is limited to a base location. With the ACP no single institution can monopolize innovation in the field of hardware development. The best way to proceed is through the development of dynamically changing architectures based on primary cost vs. performance models.
3. Current supercomputers, like old DC systems, place a cap on potential research by making it so only wealthy consumers can afford to utilize them. The new cyberinfrastructure initiative allows smaller projects to leverage greater resources without requiring large capital expenditures.

"This panel believes that the NSF has a once-in-a generation opportunity to lead the revolution in science and engineering through coordinated development and expansive use of Cyberinfrastructure."

Part III
Constellations, Globus, and the Brilliant Mr. Watt

"Grid Computing is the next big thing, worth paying attention." "What would happen if something came along that would change everything? Like the Internet, we believe that thing has come along, and that thing is the grid."

Constellations: Seeing the Bear – Ursula Major

Can you pick out the big dipper (Ursula Major constellation) in the evening sky? Typically one looks to the northern sky and finds a pot looking shape. It is very easy to see and comprehend.

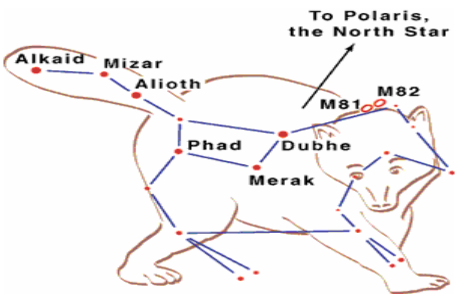

Now imagine each of these constellation points as compute clusters and the lines between them as high-speed connections. In the field of computing these systems are called grids. To better understand how these Grid/Constellations work lets look back at the APC formula for revolution. **(Distributed Computing + Cluster Computing + Data Facilities) x (Internet)/DRA[3] = Revolution**

In the future each of these resources represents a star in a constellation of linked resources. For Example, Alioth and Alkaid are high performance clusters; Dubhe is a distributed computing hub and Merak and Mizar are Data storage facilities. Dynamically linking the following resources creates ability for the birth of a true utility.

The glue that binds all these resources together is a little known program called Globus. Globus is a set of protocols and security standards which allows not only internal resources to work together more efficiently but also allows your constellation to link to other constellations for a total expansion of resources. All of an organization's tools can work seamlessly together, expanding as needed.

"The standardization of communications between heterogeneous systems created the Internet explosion. The emerging standardization for sharing resources, along with the availability of higher bandwidth, are driving a possibly equally large evolutionary step in grid computing."[4]

Say you're a marketing firm and you want to render some high-end graphics for a commercial, you only have so many computers internally, and the job would take your computers several days to complete. With a Grid, your system simply links to a set of resources and completes that job in a matter of minutes. Instead of waiting overnight or even a few days, you can find out almost immediately if any errors need fixing or any undesirable aspects of the creation need correcting.

This opens the door to the big buzzword "On Demand". All the major companies from IBM to Microsoft have on demand in their marketing and strategic direction statements. On demand is the rebirth of Tesla's practice of placing generating stations far from consumption centers, therefore maximizing efficiency.

Picture Southern California in the summer of 2001, as it was experiencing rolling black outs due to a lack of available electrical power. When Southern California finally did receive power it was supplied from places like Colorado and Oregon. The combination of multiple power generating stations coupled with grids that were able to transport it long distances allowed the electrical utility great flexibility and efficiency. The same principle is being leveraged here; grids of resources are being linked as needed to address computing problems. Linking these constellations/grids through cyberinfrastructure can allow different organizations and institutions to combine resources to the point where they can leverage up to 1000 times the computing resources available locally.

[3] DRA=Dynamic Resource Allocation

[4] -eWeek article Quoting IBM research paper

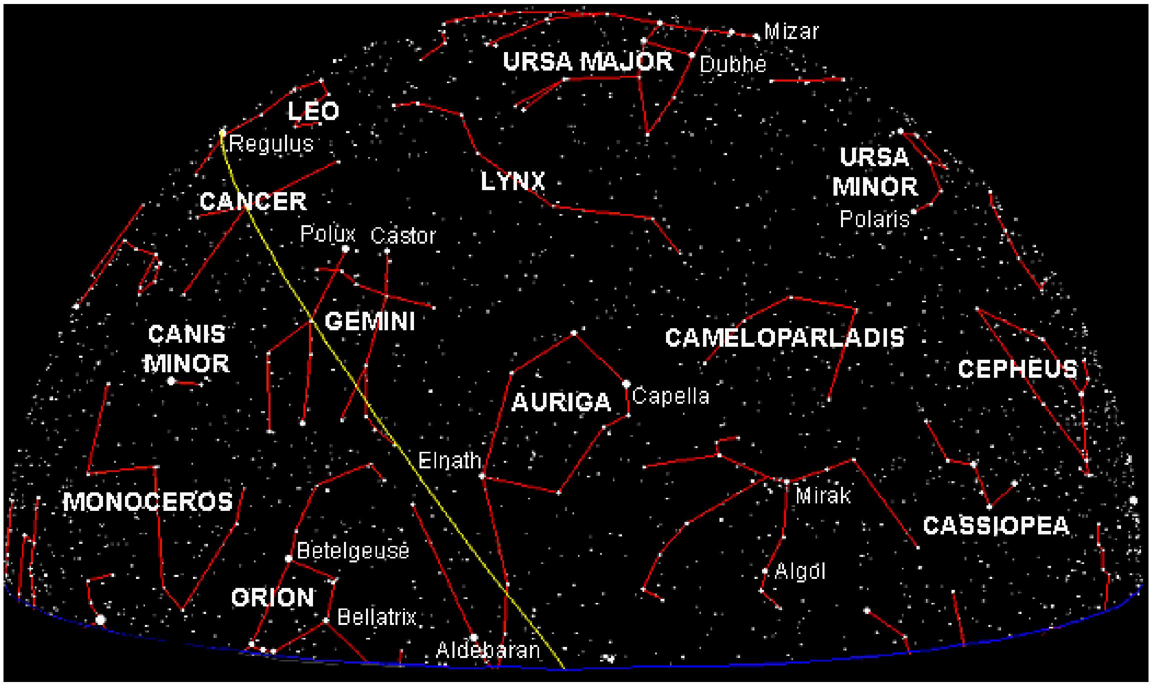

The Brilliant Mr. Watt: Monetization, Commoditization & Standardization

The other day I was speaking with a commodities exchange speculator about how the market for computational power was moving towards commodities. He asked if that was really the direction the market was heading, and if so what my opinion was on that direction? After a few days of the question meandering in the back of my mind I realized that in the development of a commodity market was not the immediate direction the market was heading. In order for it to achieve a pure utility-based model the resource must first undergo two critical steps.

1. The creation of a standard unit of measure
2. The monetization of that resource.

From this leverage point the resource can move on to a market place then an exchange finally reaching its culmination as a commodity.

Let's take a look back at the big electricity breakthrough. The brilliant Mr. Watt made the first step on the electricity path. For his work, the Scottish inventor's name is well known for the measure of electricity. Who does not recognize the "Watt or Kilowatt?" As much as this may seem trivial it is a major milestone in the development of energy. That opened the door for the monetization of this resource, which later allowed universal exchange of this resource.

The future of on demand computing also dictates a need for such a quantifiable unit of measure. In the constantly changing world of computers finding a constant unit is difficult. Crusoes, Pentiums, Athlons, and Alphas are all different types of chips working off ever changing architectures towards the end of creating computational power to drive PCs and other high-end systems.

DRI's focus on selling computational resources in a pure utility-based model requires standardization. In the case of electricity that unit is the universally understood form of watts; in the new cyberinfrastructure (on demand) model it will be Petaflops (or one quadrillion operations). This is a unit of measure any standardized computer can solve for, even with varying network technologies, processor types and system architecture variants. http://www.top500.org, a website run by several universities, utilizes the Linpack – Benchmark standard program to measure overall system performance and produce a listing of the fastest 500 computers in the world. This standard allows any system to have a quantifiable yield and monetize, or turn into cash, computational power, allowing for pure utility.

Part IV
Apples & Oranges; the Online Exchange and bring it all together.

"Digital Ribbon enables suppliers and users of computational resources, to trade, buy, and sell computational power through a many-to-many network based on an online exchange and service registry."

Apples & Oranges: Perishable fruits of the farm.

So we have a monetized standard in petaflops? Not completely. Although this is a cornerstone of the utility model one must understand that currently the market is fragmented by different architectures. If a person went to the store and was told to buy fruit they would have many choices: apples, oranges, grapes and so on. Unlike electricity, which adopted a single unchanging standard, in the field of computational power many choices have evolved. Even with a standardized unit of measure the specific need and wants may differ. Some computational solutions prefer specific ingredients. For example if you were making apple pie your recipe might call for Red Delicious instead of Crab Apples. Different sources yield different solutions; a server farm might yield Linux based solutions for rendering animation. Another company might have a recipe for rendering the same solution on a SGI based proprietary operating system. Which is like comparing apples & oranges. One rendering solution might require 4 gigs of ram and another 1 gig both on the same operating system. In this case it is a lot like choosing between green apples and red delicious apples: for some recipes this is critical sometimes it does not matter. What you need to see is that the market has variations in wants and needs.

One must also keep in mind that just like electricity these resources must be consumed or they perish. Marrying perishable resources with standardization and monetization creates a need for a marketplace.

The Online Exchange Digital Ribbons business model

As utility computing grows, the need for global market place, a place where these resources can be bought and sold in a secure environment, grows. Digital Ribbon is building that market of tomorrow, which allows the exchange of these perishable resources via an online many to many network. Our business model stands between eBay & the Electric Company. Creating an exchange that will allow both one-time transactions and varying length contracts while adding value through the regulation of Protocol & Security.

The Digital Ribbon, bringing it all together

In summary, many new technologies have been bought together, leading us into new direction of utility based computing. As utility based computing moves forward both standardization & monetization will occur. From this point markets will evolve, allowing underutilized, perishable resources to be bought and sold. Digital Ribbon stands at that junction creating a market place to manage these complex transactions while mitigating process failure through the oversight of security and system protocols. As a true utility emerges Digital Ribbon, will be a cornerstone in the forefront of the inevitable emerging market.