An [Overly] Simplified Overview of Emerging Networked Computational Technologies

Steven N. Goldstein, PhD.
Retired
Formerly, Program Director for International Networking
U.S. National Science Foundation

Introduction

I will try, in the short time available, to share with you my simplified way of comprehending the recent advances and trends in networked computational technologies. I retired from the National Science Foundation (NSF) in January, 2003, and I spent the last three years of my career there on an assignment in the Engineering Directorate where I was not in daily contact with the efforts that NSF's Computer and Information Science and Engineering Directorate (CISE) was sponsoring. Therefore, the last thing that I would want to do would be to represent myself as an official spokesperson for NSF. On the other hand, I spent most of the 1990's working to connect the academic world, that is education and research networks of other countries, to the education and research networks of the U.S. And, some of the most successful international networking projects still underway were started under my watch, so I am not unknowledgeable of the current happenings, either.

Summary

A significant and growing portion of scientific research involves distributed communities of researchers and remote resources such as supercomputers, data bases, and experimental devices such as earthquake simulators, particle accelerators, astronomical telescopes and computationally-based simulators. Increasingly capable networks are the glue that ties all this together, but a better way to visualize what is happening is to think of all the elements just mentioned as one or more integrated entities that bear a functional resemblance to the computers on your desks. Some of the key words are: Grid, middleware, lambda, optical switching, and, a real jaw-breaker, Cyberinfrastructure.

The "Grid" and associated resources

Unless you live way out in the country, you probably get your electricity, water, telephone, and maybe cable TV through a network that can be national, or even international (e.g., electricity from Canada) in scope. These networks have been referred to as "grids." When you hook into the grids, you have no idea, and probably don't really care, where your deliverables (electricity, water, programming) come from or how they wend their way to you; you just turn on the switch or the tap, and there it is! The basic idea of the computational "Grid" is the same: you plug into it and get a variety of services, and you really don't care where many of the services are being performed (of course, if you are using an experimental device like a telescope, you do indeed care about that), as long as the results that you need arrive at your desired end device (your computer, your printer, etc.).

Another way of thinking of the Grid is by analogy with the World Wide Web ("Web"). When you enter a URL (Web address), you usually get the contents of a file: data, images, animations, registration forms, etc. When you ask for something from the Grid, you will get entrée to huge databases, telescopes, simulators, accelerators, etc. This is illustrated in the graphic below. It and others in this presentation were kindly furnished by Professor Tom DeFanti (mailto:tom@uic.edu), Director of the Electronic Visualization Laboratory of the University of Illinois at Chicago (UIC). Tom has been one of the most productive recipients of awards during my watch at NSF. He continues to provide leadership for the U.S. in the emerging international arena of networked computational technologies.

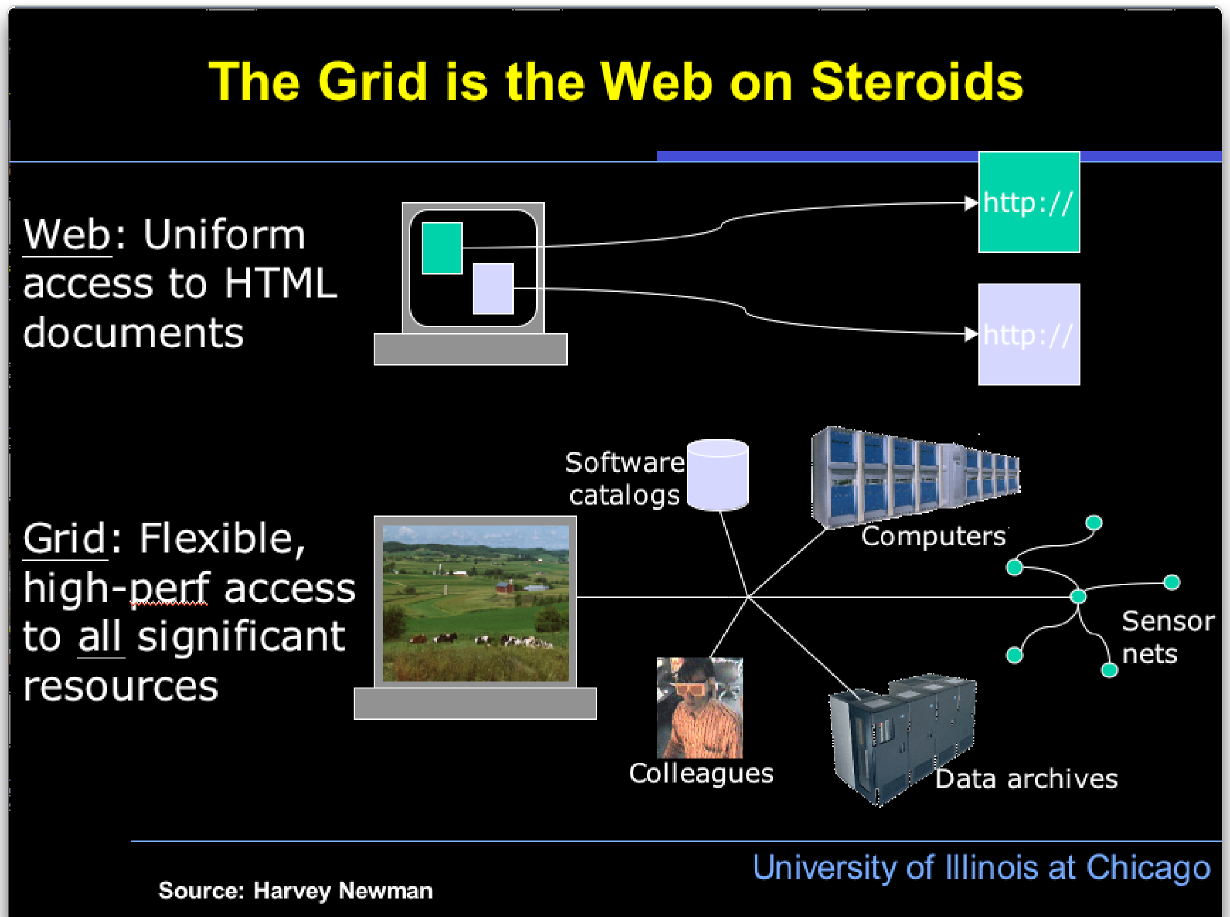

Note that the Grid is more than a network connecting all those resources. It depends on *middleware,* or a collection of software that manages things like resource reservation and allocation, authentication of users and processes, accounting, directing traffic, and security. The name "middleware' derives from its residing somewhere between the protocols that make the network work and the applications that the users employ. It is a bit like the operating system of your desktop computers, especially when you configure them for multiple users. Middleware is the glue that holds it all together and makes it all work harmoniously. There are quite a few flavors of middleware, but one of the more pervasive ones is Globus. You can read more about it at http://globus.org/.

Some examples of networking advances

People who have used the Internet over the past decade know that connection and transport "speeds" have been getting faster and faster, even with dial-up connections. As evidence, you can download big applications and animated graphics that you wouldn't even have considered asking for a decade ago unless you were on a fast university or business link. Well, that "speed" depends in large part on greater bandwidth on all the links in the chain from your computer to the computer hosting the information that you download. It means more and more lanes on the "information superhighway," including all the on-ramps. But, it also depends on the way that your information is routed. I like to describe routing in the Internet as resembling one of the old pinball machines. The information is chopped up into small packets, the pinballs. And, they are passed from specialized computers called routers to other routers, much as the pinballs hit the bumpers—except that each router has to make a decision about the next "hop," that is, which bumper to forward the packet to next; it isn't all random as in the pinball machine; there is intelligence in the system. But, it doesn't mean that some packets will not get lost, or that packets will all arrive in the order that they were sent, or that they will all get to their destination via the same route. Some specialized application--audio, for example—require that the packets all arrive in order and with a steady cadence, something that cannot be guaranteed as traffic increases in a network where everybody's packets are mixed with everybody else's (just think how difficult it is to keep a funeral procession together in heavy traffic unless there is a leapfrogging escort of motorcycle police to escort it through intersections and other traffic barriers).

Thus, in addition to huge bandwidths, sometimes the applications need specialized network protocols, either to prioritize it (which is no absolute guarantee) or to give it special reserved "lanes" where it doesn't even have to go through normal switching. Tom DeFanti explains two types of networking in the next two frames:

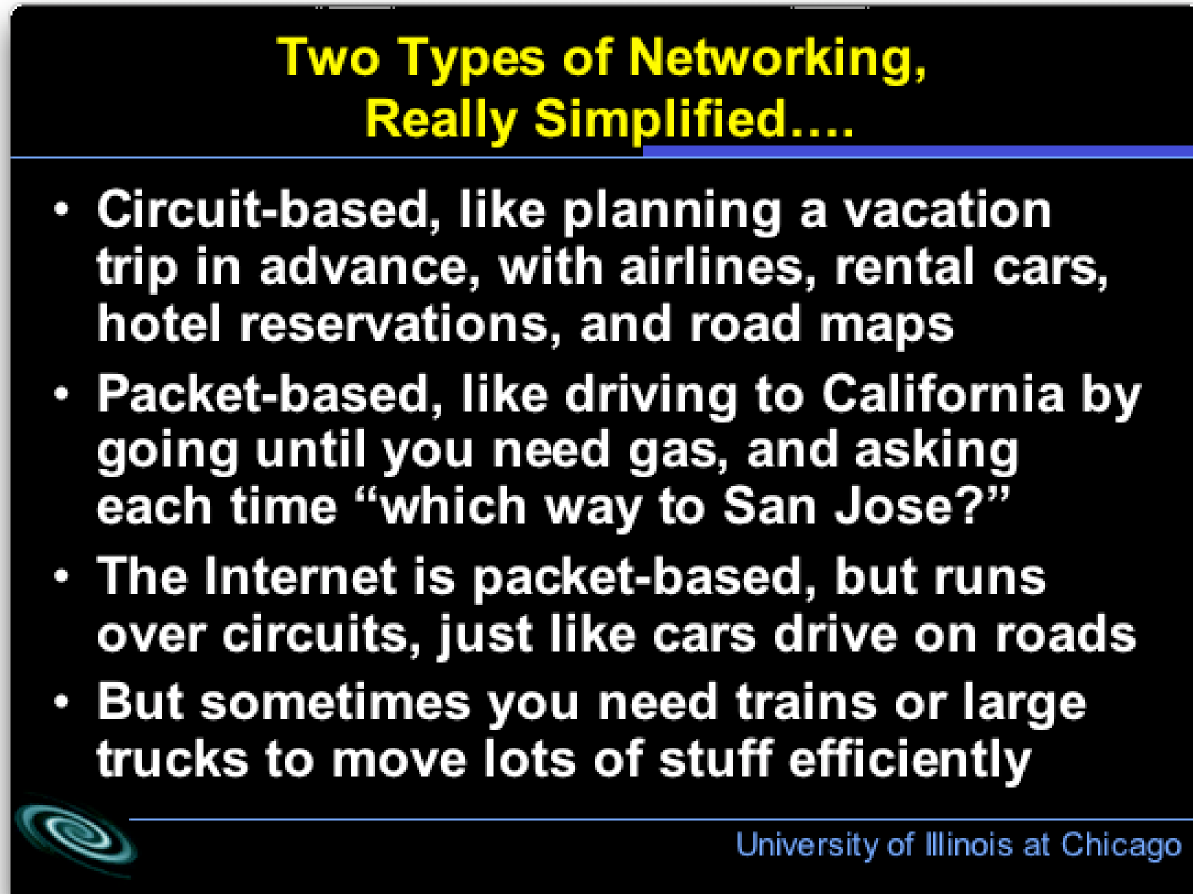

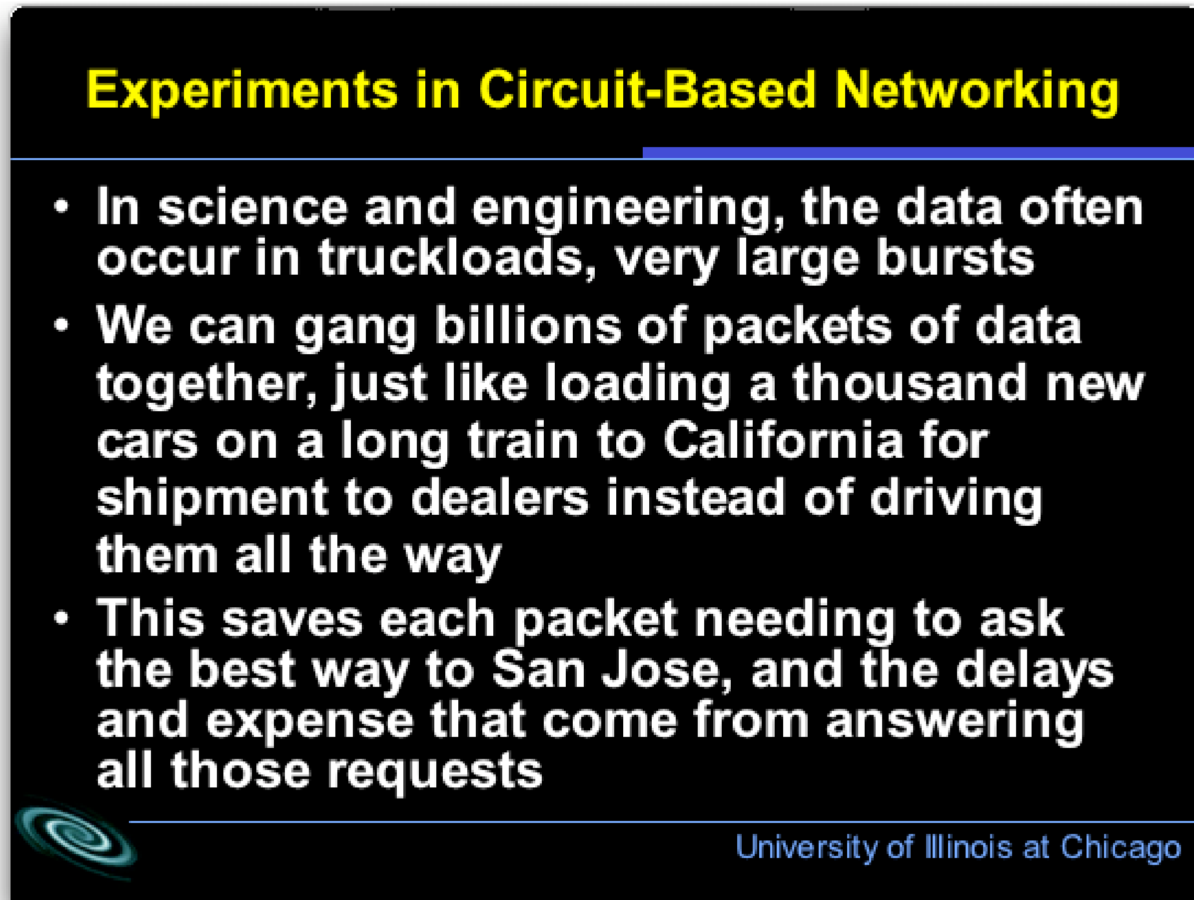

Some of the newer developments involve "hybrid" networking in which totally reserved paths can be fashioned for certain demanding applications, and its traffic passes through from origin (for example, a high-energy accelerator in Geneva, Switzerland) to the destination (for example FermiLab in Batavia, IL) without encountering routing decisions. High speed reserved links are quite rare. They are, for the most part, available only to the advanced research community and to a select few commercial outfits like the animation studio Industrial Light Management.

NSF used to support four supercomputer centers, the National Center for Supercomputer Applications (NCSA) in Urbana-Champaign among them. But, a number of budgetary and policy decisions pared the number and/or changed the mix, and one of the outcomes was that NCSA and the San Diego Supercomputer Center, plus special database (Cal Tech) and specialized modeling facilities (Argonne National Laboratory) were joined in a Gridded ensemble called the Distributed Terascale Facility (DTF), as depicted in the next frame. One key aspect of the DTF is the ganging of four 10 Gigabit/sec (Gbps, or billion bits per second) into a 40-Gbps "backbone" between Illinois and California. This is much like the main bus of your desktop computer. It is totally dedicated to the traffic that is passed among the four locations. The DTF has recently been extended to additional supercomputer facilities, but that is beyond the scope of this introductory discussion.

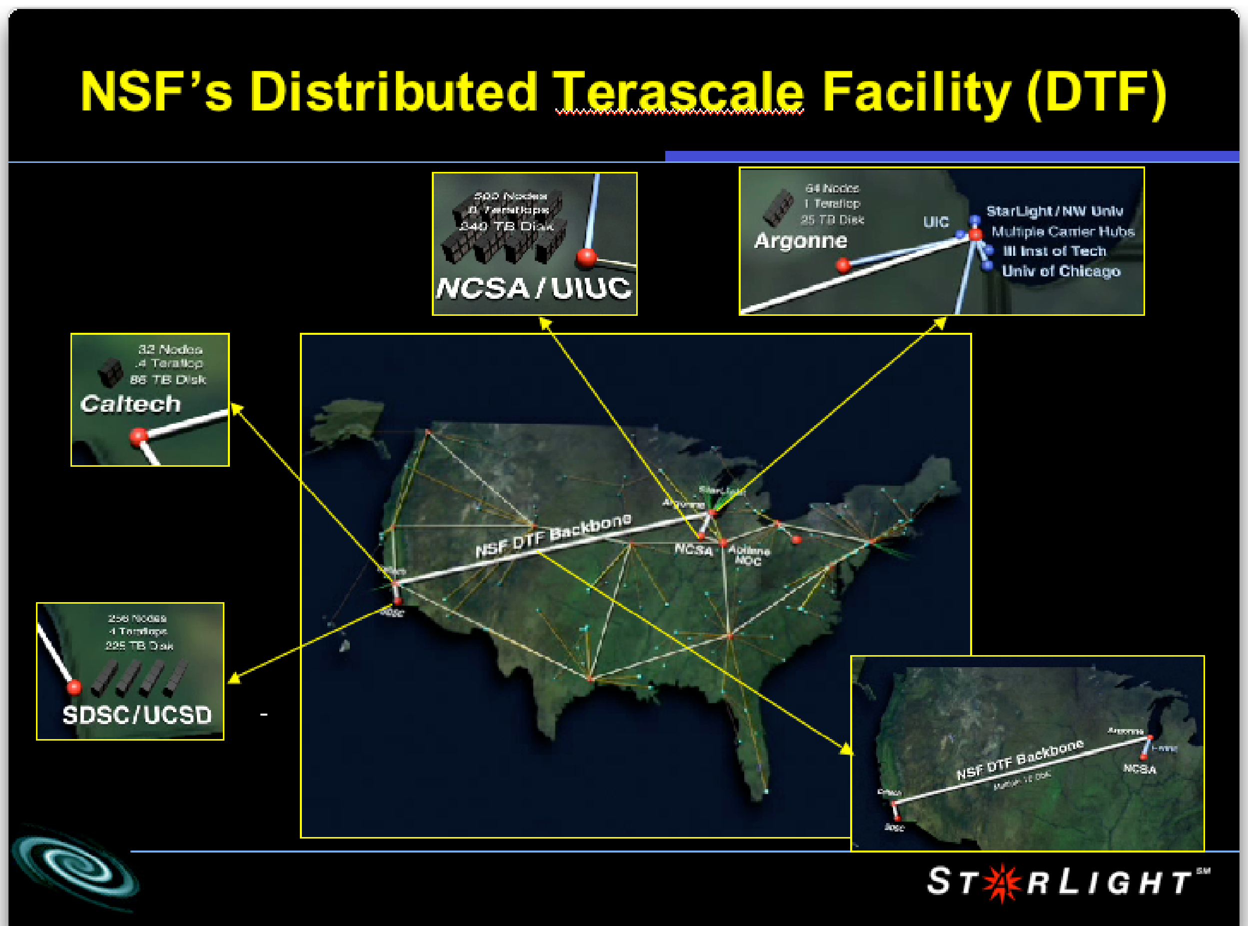

NSF's plans for extending the DTF into an "Extensible Terascale Facility" are mentioned in the excerpts from NSF's Cyberinfrastructure press at the end of this paper.

The State of Illinois has provided "dark" fiber, that is, optical glass fiber "pipes" with nothing flowing through it, and no light sources or "pumps" for university and research laboratory experimentation (next frame). An award from NSF to EVL has helped to provide the light generating and switching equipment for part of this I-Wire experiment.

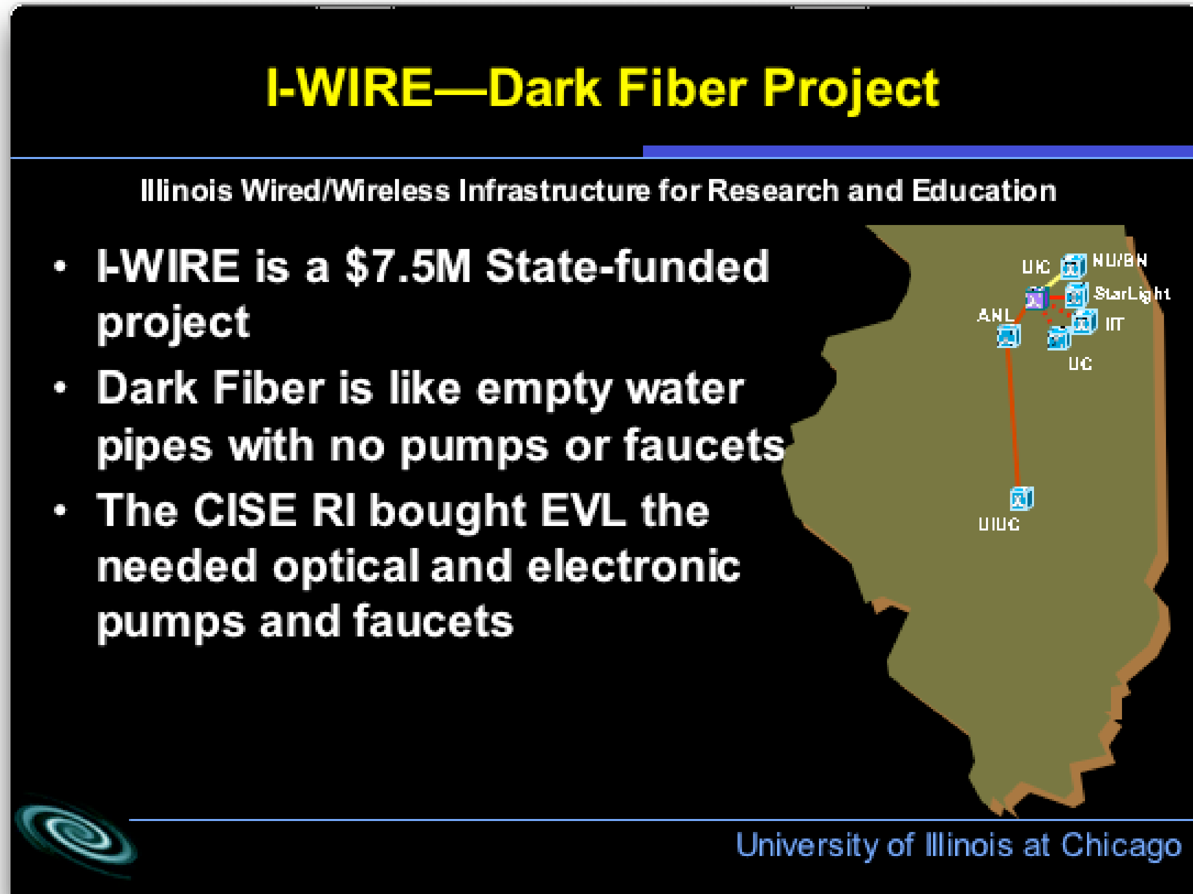

In addition, the City of Chicago has provided optical fiber to link UIC, Northwestern University, the StarLight international optical networking exchange (at 710 North Lakeshore Drive on the Northwestern Campus) and Canada's very advanced CA*net4 research network for additional networking research and development. So, in a very real sense, the area around Chicago has become a world hub in exploring new optical networking possibilities

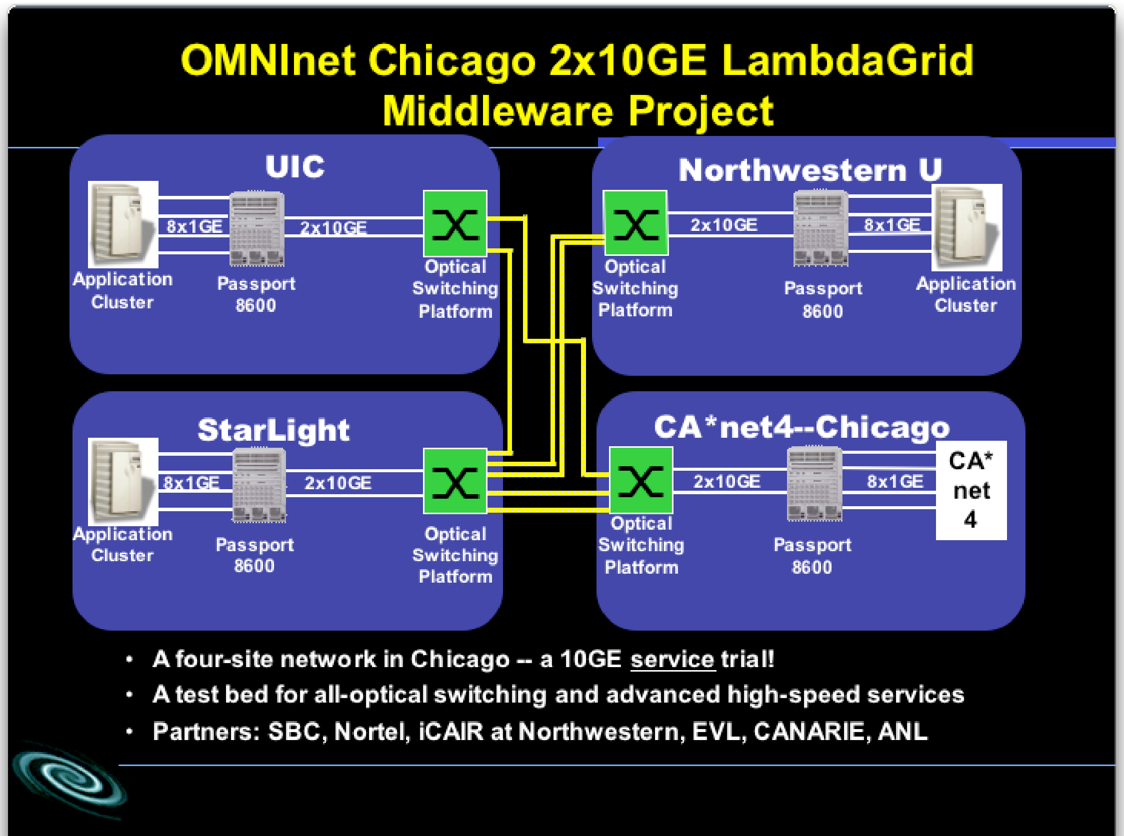

Some of the work at StarLight (http://www.startap.net/starlight) involves new *optical switching* techniques. Most all of the routing of packets or switching of circuits in the Internet is done electronically. Even when the links are optical, sending laser light down long fiberoptic paths, the optical signals have to be converted to electronic signals for routing and switching, and then back to optical again for further transmission on the fiberoptic paths. With optical switching, there is no intermediate conversion to electronics; all of the routing and switching is done in the optical domain with new devices such as very small multi-faceted mirrors. Again, this is very advanced stuff, and industry is an active participant in furnishing their developmental equipment for trials.

As to the vocabulary building: the Greek letter *lambda*, λ, is the standard symbol in physics for the wavelength of a color of light. Over the last decade technologists learned that they could increase the amount of information that could be sent along an optical fiber by shining multiple colors, or wavelengths of light. Now it is not uncommon to employ 60 or 100 distinct wavelengths, or lambdas, to increase the carrying capacity by a factor of 60 or 100, each lambda's being independent of the others and therefore acting as if it were a separate pipe. So, optical switching switches individual lambdas, each of which can carry 2.4 or 10 Gbps of data.

Together with his colleagues in The Netherlands, Kees Neggers, and in Canada, Bill St. Arnaud, Tom DeFanti started a trans-Atlantic collaboration with a growing number of 10 Gbps lambdas. (There is a bit of terminology-stretching going on, too: as a bootstrapping step along the way to true and completely optical switching, some of the 2.4 or 10 Gbps lambdas were actually subdivided in the electronic domain to Gigabit Ethernet speeds of 1 Gbps or less and switched electronically. For convenience sake, they called these subdivided links "lambdas" as well. But, rest assured, the move is in the direction of pure optical switching with real honest-to-goodness optical wavelengths.) Other national research network entities have joined in the project, which is now called TransLight (http://www.startap.net/translight/ -- next frame).

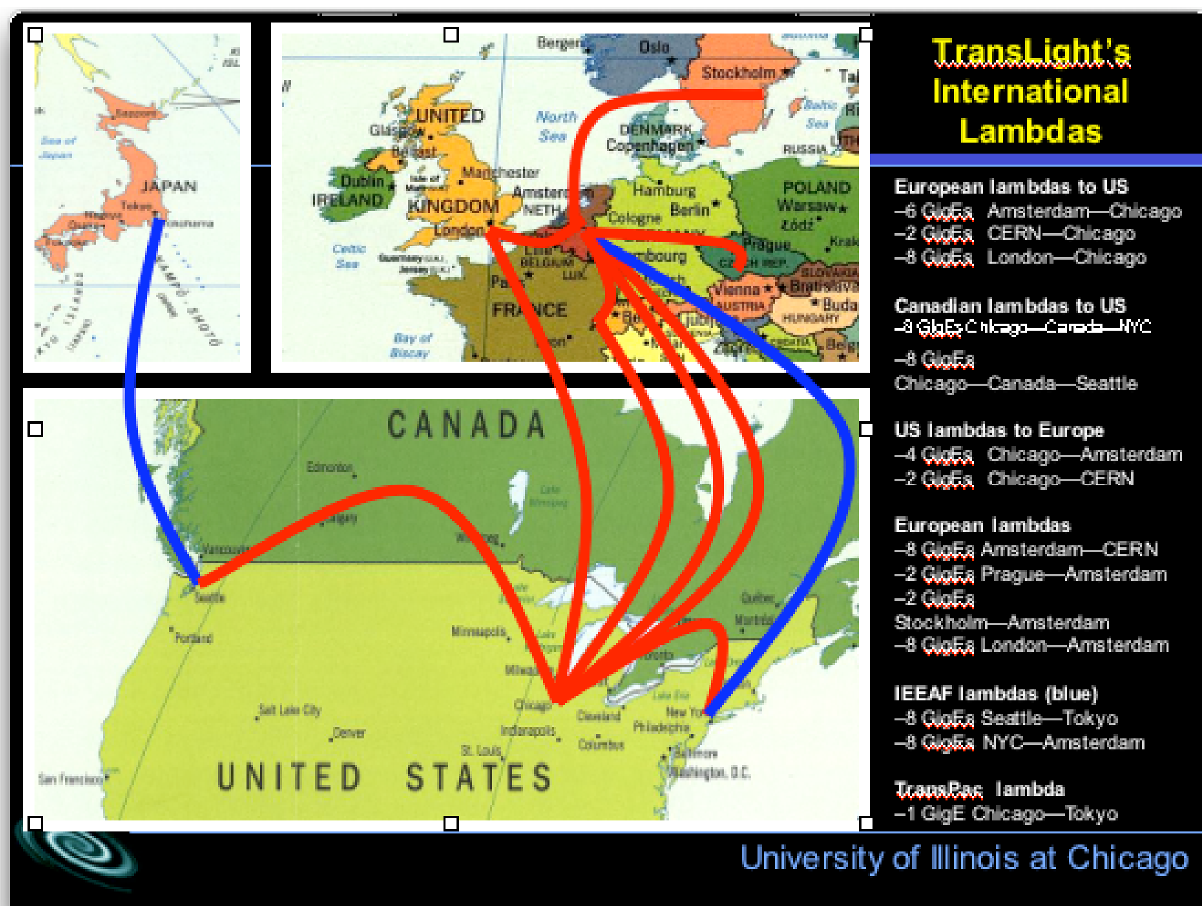

The truly networked computer, OptiPuter

As a final example, I briefly note that Larry Smarr, NCSA's former Director, and Tom DeFanti have put together the OptiPuter project which is described in the next two frames.

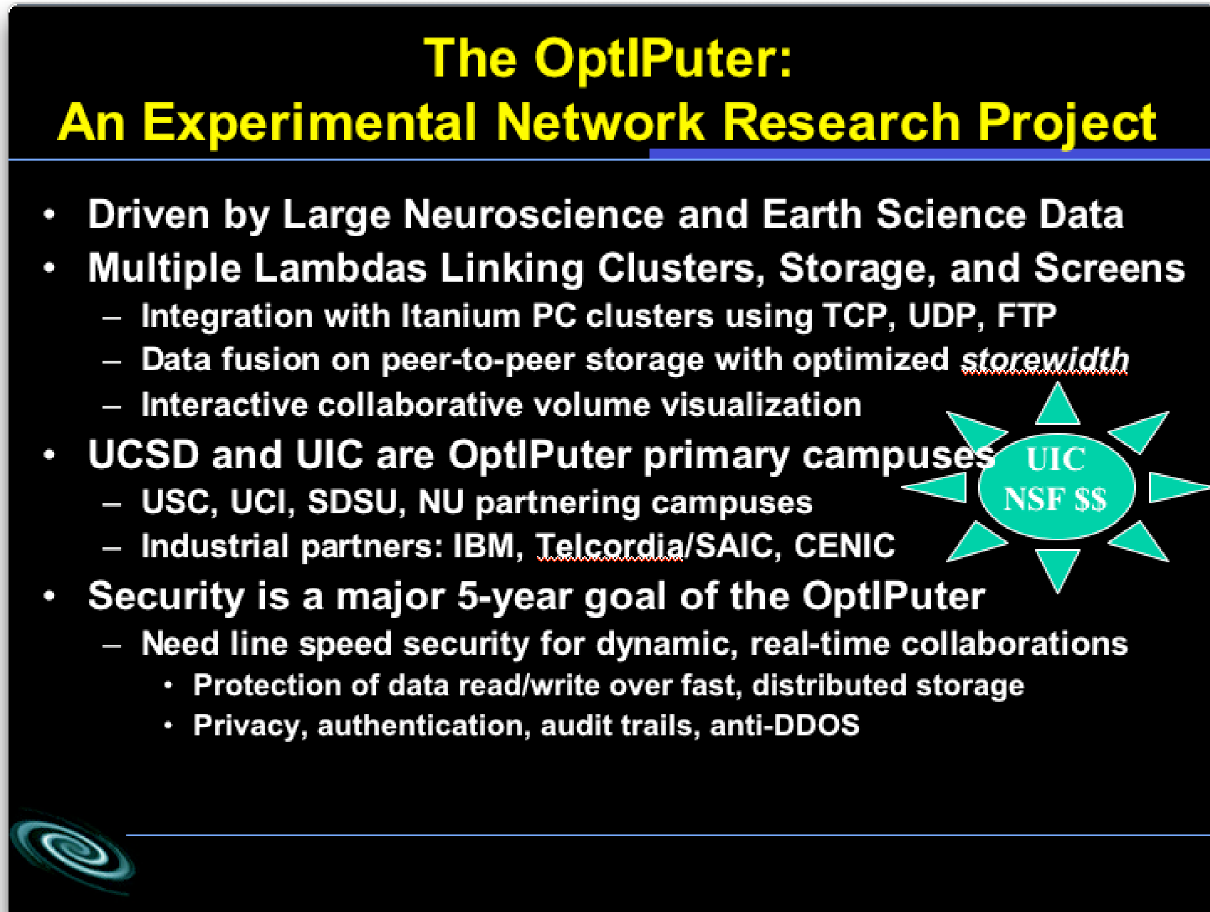

Basically, the OptiPuter is like a huge desktop computer whose various pieces could, in principle, be distributed all over the world. One of the objectives will be to "tune" all the parts so that they will work smoothly together. To quote from the announcement:

> "The OptIPuter can be seen as a "virtual" parallel computer in which the individual "processors" are widely distributed clusters; the backbone network is provided by IP delivered over multiple dedicated lambdas (each 1-10 Gbps); and, the "mass storage systems" are large distributed scientific data repositories, fed by scientific instruments as near-real-time peripheral devices."

(http://www.startap.net/starlight/PUBLICATIONS/news-optiputerBandwidth03.html)

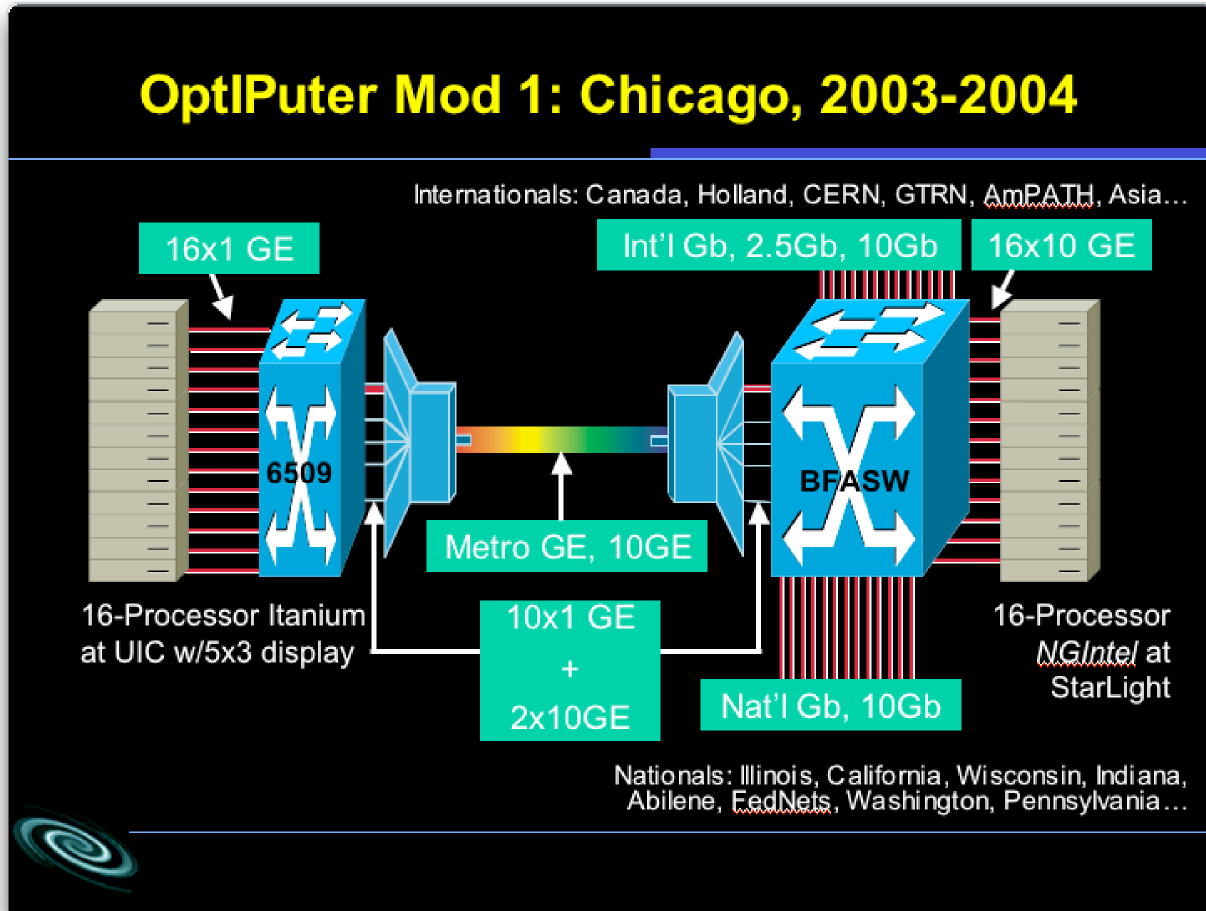

Cyberinfrastructure

OK, now for the jawbreaker. We have hardly scratched the surface of the various parts of the Grid, actually the many Grids that serve a multiplicity of scientific and business communities, let alone the scientific communities and their application needs that depend on seamlessly functioning Grids. But, put them all together, and they spell "Cyberinfrastructure." The former NSF Assistant Director for CISE coined the term several years ago. But, it is a bit like the proverbial elephant with blind men examining and reporting on it; there are several slightly different definitiions. You can find the official Report of the NSF Blue Ribbon Panel on Cyberinfrastructure at http://www.communitytechnology.org/nsf_ci_report/. Just a short quote from the Executive Summary illustrates the totality of the concept:

> **The emerging vision is to use cyberinfrastructure to build more ubiquitous, comprehensive digital environments that become interactive and functionally complete for research communities in terms of people, data, information, tools, and instruments and that operate at unprecedented levels of computational, storage, and data transfer capacity.**

Overview of Emerging Networked Computational Technologies

On October 28, 2003, NSF announced "continuing steps to ENHANCe Cyberinfrastructure" http://cise.nsf.gov/news/cybr/cybr2.htm :

> …During the past summer, NSF solicited advice and input from the academic community through two workshops and a town hall meeting to discuss management and models for cyberinfrastructure. Internally, the agency convened a Cyberinfrastructure Working Group to explore challenges and opportunities in all science and engineering fields. With guidance from these and other sources, NSF has outlined the following steps in the ongoing transition.
>
> As previously announced, the Partnerships for Advanced Computational Infrastructure (PACI) have been extended through the end of FY 2004. During this period, both PACI lead sites - the National Center for Supercomputing Applications (NCSA) and the San Diego Supercomputer Center (SDSC) will deploy significant technology upgrades, almost doubling the high-end computing resources that NSF makes available to the nation's scientists and engineers. NSF plans include revised agreements with NCSA and SDSC to ensure the continuing provision of high-end computing resources and related services to the national community through the end of FY 2007. In addition, SDSC and NCSA will work in partnership with NSF and the science and engineering community at large to define emerging cyberinfrastructure opportunities to advance all fields. These and other community activities will inform NSF's development of future cyberinfrastructure-enhancing competitions.
>
> Complementing the cyberinfrastructure resources and services provided by NCSA and SDSC, the Extensible Terascale Facility (ETF) - which is on track to be commissioned October 1, 2004 - will demonstrate the potential of revolutionary grid computing approaches to advance science and engineering research and education. Additional ETF upgrades are being planned for FY 2004, where plans include new capability computing for the Terascale Computing System at the Pittsburgh Supercomputing Center (PSC). This represents the final stage of the ETF's construction phase. Support for the management and operations of ETF-enabled cyberinfrastructure will be provided beginning in FY 2005 and extending through FY 2009.
>
> To ensure that all science and engineering communities are prepared to inform the development of and effectively utilize the broad, evolving cyberinfrastructure, NSF plans to hold an open competition during FY 2004 that will ultimately support a comprehensive set of education, training and outreach awards. This competition will build on the work of the successful PACI Education, Outreach and Training (EOT) and other activities. …

So, there you have it!