The COOK Report on Internet Protocol
Technology, Economics, and Policy
© Cook Network Consultants

ICT and E-Science as an Innovation Platform in The Netherlands

A National Research and Innovation Network What Can the US Learn from Dutch Experience?

Volume XVII, No. 11
February 2009
ISSN 1071 - 6327

"The dogmas of the quiet past are inadequate to the stormy present…As our case is new, so we must think anew and act anew." **Abraham Lincoln**

By means of an examination of research networks in Holland, this issue presents some ideas for ways in which an American National Research, Education and Innovation Network could be structured.

For the first time in more than a generation, the model of unregulated speculative financial capital has shown its bankruptcy, the ability of government to encourage the coordinated use of society's resources in the public interest should become a focal point of our political life.

By encourage I don't mean dictate, but rather to try to act on behalf of agreed upon basic principles that wherever possible are carried out by decentralized groups.

What can be done privately should be. But government must exert oversight and insist on transparency. We can only hope that the new Administration will begin to explore these and many other new ideas.

Before Resource Use Careful Deliberation and Coordinated Planning Needed

In the area of networks as an integral part of national social, economic, and research infrastructure The Netherlands is now a world leader. The Dutch are building a national, largely open, fiber infrastructure. As we demonstrated in our January 2009 issue, The Netherlands has a pragmatic way of finding out "what works" and then just doing it. In this issue we examine in detail their current research and innovation network infrastructure.

In SURFnet6 the Dutch have likely the most innovative optical NREN in the world. To leverage the capabilities of SURFnet6, in 2005 they started ICTRegie – the Netherlands ICT Research and Innovation Authority. The mission of the Netherlands ICT Research and Innovation Authority www.ictregie.nl is to enhance the innovative strength of the Netherlands

On the Inside

Pioneering LightPath Networks for Global Science Collaboration
Contents p. 65

Please read Explanatory Note page 64

by means of information and communications technology. Two key objectives derive from the mission statement. First to introduce unity and consistency to the strategic direction of the ICT research and innovation by means of the development of a national strategy that enjoys broad support and, second, to ensure ongoing strengthening and appropriate dynamism of a Dutch ICT knowledge infrastructure geared towards high social and economic yield. ICTRegie has a Strategic Plan 2005-2010 found at http://www.ictregie.nl/index.php?pageId=6&l=en&personeelId=&pubId=15

This Plan is designed to establish a framework for thinking, discussion and coordinated action to ensure effective use of economic capital invested in the goals of the organization. It strikes me as somewhat similar to the Japanese ICT plan of 2000. Other plans from other countries could be identified. What troubles me is that this approach in the United States would be likely to be criticized as state planning. Of course the five year plan of the 1930s remains a joke.

What we don't seem to have considered and I suggest that we should consider is whether a plan for coordinated use of resources in a framework for such can be far more productive than throwing huge amounts of money in an uncoordinated way and at a wide array of problems in the hope for some results. The modern world is too complex and capital in too short supply for the continuation of this profligate way of governing.

Here is a world view we need to emulate:

The four pillars of ICTRegie's strategy

1. The Netherlands ICT Research and Innovation Authority (ICTRegie) was founded with a view to strengthening the country's ICT knowledge infrastructure (in both focus and mass) and to match supply to demand in terms of knowledge and applications. The overall aim is to enhance the innovative ability of the Netherlands (impact).

2. Within the complex field of ICT, its many research disciplines on the supply side, countless application domains on the demand side and a large number of organizations throughout, a purely top-down style of direction is unlikely to be effective.

3. Three accordingly, ICTRegie *mobilizes* stakeholders on both the supply and demand side, encouraging them to join each other in thinking about opportunities for innovation. It *challenges* them to inspire each other to arrive at promising innovations using ICT, *encourages* them to seize the opportunities for innovation, and brings those opportunities together within a national vision.

4. Making the best possible use of self-organization abilities of the field, ICTRegie develops new instruments and programs in selected fields.

E-Science

The remainder of this issue focuses on E-science as developed under the leadership of Kees Neggers and Cees de Laat at SURFnet in conjunction with others in the United States (including Tom DeFanti, Larry Smarr, Maxine Browne, Joel Mambretti),Japan, Korea, China and Europe.

Kees Neggers sent me the December 11, 2008 ICTRegie report called "Towards a Competitive ICT Infrastructure for Scientific Research in the Netherlands". This report can be found at: http://www.ictregie.nl/publicaties/nl_08-NROI-258_Advies_ICT_infrastructuur_vdef.pdf

The report evaluates and builds on both SURFnet's and the e-Science Virtual Laboratory project's innovation approach. I refer readers to the conclusions in sections 6.8, 6.9 and 7.0 found on pages 59 and 60 of this issue. Trying to place my emerging understanding of e-science in context I asked Kees by e-mail

"Let me test my understanding: There is

1. theoretical science
2. applied science (building something from theory).
3. computational science building a computer model to see if the application works and finally
4. e-science

E-science is not really an independent 4th branch of science. In reality it is more like a platform or foundation with which to enable the efforts of scientists at all three levels to support the economic foundations of their society.

How to describe it in more detail? Using light path networks - OptiPortals - data storage, grids for real time collaboration of people. and instruments designing projects around the world - it seems to be the internet based emerging global platform for all science, teaching and investigation globally?

Neggers: I agree, actually it is a push to a "better internet" too.

The ICTRegie report defines e-science in chapter 6.5 as "a development to bridge the-gap between scientist in application domains and the development of ICT." In presenting the report, the director of ICTRegie called it enhanced-science. In the preparation of the report some suggested to avoid the word e-science completely because of its unclear meaning.

COOK Report: But if theoretical science ever advances from theory to practical product, if must depend on these e-science platforms?

Neggers: Very likely - yes. We see a growing demand emerging in all disciplines. The Dutch Roadmap Committee for ESFRI already put more emphasis on alpha and gamma science and as a result on data versus computation. This will not replace the need for computation of course. It shows that infrastructure like networks, computation and storage, plus the required software and know how will be essential tools for researchers in all disciplines soon.

And that's why it should be made available, at high quality, in its own right. To avoid duplication and/or delay in realizing it.

Adopting a New National Infrastructure

How Fiber Became the Basic Platform for Education, Commerce, Technology Transfer, and Research in the Netherlands

Editor's Note: I would argue that The Netherlands has a more complete and pervasive fiber optic infrastructure than any country in the world. On the citizen side fiber to the home and apartment and on the research education and enterprise side not only fiber but also light path networks to 160 universities, research institutions and enterprises. This issue revisits SURFnet for the first time since the March-April issue of 2005.

In this opening interview, done on Monday November 17th at Supercomputing 08 in Austin Texas with SURFnet Director Kees Neggers, we look at the events that enabled the Netherlands to become both the optical and IP focal point for Europe.

Sparking the Fiber Revolution

COOK Report: How did the fiberization of the Netherlands begin?

Neggers: In the 90s the Netherlands liberalized its telecommunications law. In the new law, which came in effect on December 15, 1998, all land owners, both public and private, had to tolerate the digging of fiber by all public network operators. Public operators did not have to pay for use of right of way. Until than, only the incumbent operator, KPN, had these rights. Also operators of non-public networks were allowed to lay there own fiber, but they had to make an agreement for this with the land owners first. City Governments were given an official coordinating role to streamline the diggings. Another important aspect of the new telecom law is that when you make use of your right of way, you have to allow sharing of the digging. In other words if you begin a project you must announce what you plan to do and anyone who also want to lay fiber as part of that excavation is entitled to do so as long as that entity shares the cost of the excavation appropriately.

This all meant that many more were able to play and,

especially in the end of the 90s, a lot of fiber was laid. Particularly in Amsterdam. See picture from an old presentation at that time above. All of this was put in before the breaking of the .com bubble. All this excavation included a large number of empty ducts that could easily receive additional fiber.

COOK Report: If there are disputes, what happens?

Neggers: The cities have to deal with disputes for fiber digging activities in all public grounds. OPTA, the Dutch telecom regulator, was charged to deal with disputes for digging in private territories. In both cases of course also the standard legal appeal procedures are available.

The Significance of AMS-IX

Neggers: The Amsterdam Internet Exchange (AMS-IX) began in the early 90s when the Internet was primarily a research network. But early on commercial players were allowed to interconnect and a very important decision was that the commercial carriers were allowed to interconnect with the early research networks on a neutral basis. This meant that, very early in the game the research networks and the commercial networks in the Netherlands were very well interconnected.

De Laat: The first two locations for the Amsterdam Internet exchange were actually both within the Science Park.

Neggers: We started in two places - at SARA, that exchange was run by us and at NIKHEF, a Dutch High Energy Physics site. We designed an exchange architecture with a high priority on resiliency - with a policy that encouraged people to connect in both locations. After a few years we formed an association to create a level playing field for all participants. In order to be connected to the exchange you simply had to join the association as a member and abide by its rules. **AMS-IX is still a not-for-profit association and is not controlled by any single member.**

Because a neutral exchange existed in Amsterdam very early on, as the Internet became more important, it served as a magnet that enticed operators to bring in connections. To take advantage of the conductivity, all they needed to do was to bring their fiber duct into the AMS-IX location which then served as a catalyst for this continued growth.

And when we had SURFnet5 up and running in 2001 with 10 Gb Lambda's connected to Cisco GSRs, we wanted our customers to be able to connect at 1 Gb or 10 Gb speeds as well. We had asked the providers for Gigabit Ethernet connections, but operators were unable to provide these in 2001. Of course, to deliver the gigabit Ethernet service, the operator would have to bring in fiber. Realizing this we decided why not just ask for the fiber! We then asked the operators to give us the fiber pairs we needed for our locations. Again, they refused to deliver.

We also realized that individual trenches dug to single locations for every SURFnet5 member would not make economic sense. Therefore we planned fiber rings through the areas where we had clusters of institutions connected to SURFnet5. We did this in conjunction with city governments which joined in the procurement so that they could also connect their institutions to the fiber ring.

The operators did not even respond to our tender which of course left us no choice but to go to the specialized fiber installation engineering groups to which the operators had outsourced their own building before the crash caused them to stop further capital investment.

Due to the dot com crash, the operators were no longer investing in capital expenditures and consequently, the companies to which they had outsourced the building of their own fiber networks were standing idle and were in need of work which we were glad to give them. Because of these economic conditions the first ring with a length of 40km with 20 locations and 96 fibers did cost only €1 million. The return on investment for the entire ring was less than four years based on 2 Mb per second leased line prices for these 20 locations.

Consequently this was a no-brainer. The carriers soon discovered that they would have competing fiber rings over every square inch of their territory if they did not make their own fiber more available, which they then did. This was a tipping point that actually opened a dark fiber market for everyone in the Netherlands.

Now, had they provided those gigabit connections, we might have become hooked on them and not have started to build them for our own use. We never thought that having to dig ourselves would become so attractive. Of course, had it not been for the "dot com" crash and things were booming, then the technical civil engineering companies who are doing the fiber laying work would probably not have been available to us. And another nice thing was that these companies were the ones who originally had done the digging for KPN, AT&T, Level 3 and other carriers so they knew very much what they were doing and were quite capable of providing good work for us.

It was also fortunate for us that the operators were no longer like the old PTT because most likely they would then still have completely owned and controlled all engineering installation operations.

COOK Report: Gradually then, by means of this strategy, you had all the fiber procured that you needed?

Neggers: Yes, and by about 2005 we had connected all 160 SURFnet member institutions directly by fiber. This was local loop dark fiber with a 15 year IRU owned directly by us. We completely solved the last mile problem by means of these local loops.

But then, when the contract with British Telecom ended and the SURFnet5 backbone, based on BT 10G lamdas, had to be replaced, we said; why not put a little bit of extra fiber in the backbone and go all optical? This is what became SURFnet6. In effect we had all last mile connections in dark fiber and only needed some additional dark fiber in the backbone. This was the difference between us and other NRENs that are going hybrid now. These NRENs all start with the backbone and then still suffer from the last mile problem.

Because we have an all optical network, we were able to design it so that we had optical nodes at our largest institutions. Consequently, if we needed to bring in additional connections, it was easily doable and all-optical as well. The cost of SURFnet6 was similar to the cost of SURFnet5. SURFnet5 was a 10G backbone with 15 big GSR routers. With the all-optical approach we reduced the number of router backbone sites from 15 to two. We used two for resilience. Both function all the time and everyone is connected to both. We also have a border router at each of our two backbone sites to separate external from internal routed traffic. So the whole SURFnet6 IP network is based on 4 routers in two sites.

COOK Report: It was this idea that was the germination of your hybrid network approach?

The Germination of the Hybrid Network Approach in 2001

Neggers: No, the seed for hybrid networks was in 2001 when we had these 10 Gb lambdas from British Telecom and realized we could do a lot more with them, specifically for demanding users like high-energy physics and astronomy. We also realized early on that the research world is an international community for which we had to develop lambda networking on an international scale from day one to avoid inter-working problems in the future. Consequently in 2001 we first ordered a 2.5 gig from Amsterdam to StarLight, the US research networking exchange point in Chicago, in order to be able to explore these options with Tom DeFanti and Joe Mambretti.

By iGrid2002, which was in Amsterdam in September 2002 we already had two 10 Gb lambdas between Amsterdam and the US as well as the 2.5 gig wave. But nobody could really use them at that time. All of the TCP/IP stacks and the workstation Ethernet cards crumbled. They had all been testing the stacks in the lab where they worked but suddenly with these long distances they failed. However, everyone realized the potential and the lambdas themselves were there to stay.

In 2001 we also planned a dark fiber to the Dwingeloo location where our astronomers are and where the European JIVE correlator is. (You see that some time ago the Netherlands was rather interested in navigation and therefore we invented all sorts of tools like telescopes and still astronomy remains a significant research field in the Netherlands.)

Global Crossing, which was also having tough times in 2001, had a cable from Amsterdam to Hamburg that went within 40 km of Dwingeloo. Global Crossing was prepared to take a fiber pair from Amsterdam to Hamburg and open it up at the point nearest to Dwingeloo, if we'd dig a 40 km trench from that location to the Dwingeloo Research Center.

Adoption of Lightpaths - Not only for Science but also for ICT Departments

COOK Report: So where have you taken your SURFnet architecture in the last three years?

Neggers; The SURFnet6 architecture worked out pretty much as we intended. Nortel won the procurement as you know. We were the first customer for their new Common Photonic Layer (CPL) equipment and we are still very happy with that. The network is totally optical all over the country. We installed their Avici routers but unfortunately they withdrew from the market. Nevertheless with our collapsed backbone design we have only two such routers connecting all our customers and we were able to make a transition to new Juniper routers that was transparent to our users. We installed them back to back with Avici's and every week moved a number of 10 Gb connections from one platform to another. We did so in a maintenance window so that customers did not notice the change.

In the early 2000s, when we were experimenting with light paths from the Netherlands to Chicago, there was still a lot of hesitation in Europe. I remember that I had a presentation at the NORDUnet conference in April 2002 titled "Research networking – The next Phase" where I said this is the future this is where we have to go. Even NORDUnet was not ready for this this at that time. Most people thought that 10 Gb IP, which was already running in SURFnet5, would be the future and bring plenty enough for everyone. But we went on this path anyway because we were convinced that we had no choice. Our astronomers were already telling us that this was their only future: by having telescopes connected at high bandwidth to the JIVE correlator they could do real time correlations. So we also connected our 16 linear array antenna radiotelescope in Westerbork to JIVE in Dwingeloo.

Consequently, when we started developing ideas for SURFnet6 in 2003, we had this footprint of all-fiber-connected institutions and it was not a to big expense for us to do an all optical hybrid network to provide both optical and IP services. For us it was a no-brainer. But we realize that for everyone else who has not yet built the fiber infrastructure to do this it will still be a major investment. But today there is no escape for anyone involved in research networking, and if you do not do this, you may just as well walk away.

When we built this hybrid network motivated by our big science users, we found to our surprise that the biggest uptake in the Netherlands was not with the researchers but rather with the ICT departments. The reason was it is the ICT departments that are responsible for providing network services through to remote campuses for e-learning. They had problems doing this by means of a routed IP network. They saw that they could order lightpaths from us, and use them to extend their local area network, and their problem was solved.

COOK Report: This is very interesting because it must mean that there are potential lightpath applications in connecting remote corporate campuses with each other via optical private networks?

Neggers; Absolutely. If you can have a centralized ICT department concentrating your services at one or two (for resilience) places then, as a result, you can have the same quality of service to all your researchers or, as the case may be, to all your remote corporate campuses. Now the people producing software for these services to remote campuses will find that they need to take care of the latency problems.

For example with iGrid2002 we found out that they had not thought about the problems of remote users scattered over a very large wide-area network. With ordinary speeds and TCP a dropped packet didn't matter. But with gigabit or multiple gigabits light waves running on TCP if the packet was dropped the stack would crash. We also discovered that login protocols for access to remote servers not always know how to deal with WAN latencies.

In the Netherlands, when we started SURFnet6 we estimated that by the end of the GigaPort project in December 2008 most of the big universities and a few large research institutions would be connected to the lightpath network. Now however we already have about one third of our connected institutions using lightpaths. and most of these institutions are using lightpaths to create Optical Private Networks for e-learning and administrative services. On a domestic basis the uptake is ICT departments. While on an international basis the uptake is all researchers.

Internationally we have GLIF, the Global Lambda Integrated Facility, collaboration where we share lambdas and coordinate the scheduling of lightpaths for experiments. This is very successful. Right now many researchers are making serious use of lambdas internationally where IP routed traffic at such high speeds and quality is really not feasible. Payment problems for the researchers have been avoided so far because we have had a lot of donated or shared lambdas for international experiments in use.

But nationally, even in the Netherlands, there is still a lot of help needed to make the facilities available to and used by the applications people, that is corporate scientists or university researchers. Researchers are faced with local loop problems inside the campuses, needs for reallocation of ICT budgets, and last but not least, restructuring of their research projects to be able to take advantage of the new facilities. All this takes time and effort. We therefore have to do much more to stimulate the uptake of the optical network by the research community.

Nevertheless the fact that about one third of our customer base in two years time is actually using lightpaths, we rate as quite a success

COOK Report: Are you in a position where you could provide training to people outside of the Netherlands who wanted to embrace this technology?

Neggers: As you have seen we, and the research networking community in general, are quite open about what we do and how. This is also quite essential. Otherwise it would not be possible to create a transparent multi-domain lightpath service to our international user base.

GLIF of course is an excellent vehicle for this and I would recommend that anyone interested should join this community. Also users are welcome to join the GLIF Research and Application Workgroup. SURFnet is organizing seminars and training for our connected institutions. But we are not in the business of doing this commercially because we believe that there are enough other people in the Netherlands who can do it certainly as good as we can and perhaps better. In other words doing this is not a part of our core business.

Dutch to American Analogies and the E-Science Virtual Laboratory

When you consider what you might like to try in the United States you have to realize that the situation in the Netherlands is very much different than in the US. NL is very much smaller. In the US terms the Netherlands is very much like a regional US network. But remember that in addition to the concept of the Netherlands as a regional there is also one more huge difference. All our research institutions are connected by fiber. There is no last mile problem towards the institute anymore.

If anyone wants to connect to a Dutch university or research Institute, the only thing they have to do is bring a lightpath into NetherLight, the Amsterdam Glif Open Lightpath Exchange. Consequently **in the Netherlands a broker with knowledge of industry and the research area could just begin to function as an application broker because the network is already there. To do the brokering all you have to do is make sure the parties, or their providers, are connected to NetherLight.**

COOK Report: To what extent and do you have good cooperation between university and private industry where the desired end is technology transfer?

Neggers: Cees de Laat can now take over because he is part of the Dutch Virtual Laboratory for E-science (VL-e) project, which is the other project that, in conjunction with GigaPort Next Generation Network (the project that developed SURFnet6's optical network), was funded out of the government supported so called Bsik program. When we started the GigaPort project in 2003 we joined forces with the University of Amsterdam's VL-e project. We submitted a largely complementary proposal and said, if we are both funded Cees his group's networking activities will be in the GigaPort project. But he also had some non-networking responsibilities in the virtual laboratory e-science project and this is the area which your question now directly touches upon.

Cees de Laat: The Virtual Laboratory for E-science (VL-e) has university partners and industry partners. On the industry side for example there was Unilever, IBM, and Phillips, the Dutch electronics company that operates a large scientific research park in Eindhoven. Now the BIG-GRID project, a sister project of VL-e, has build a considerable grid. Part of this infrastructure is installed in Phillips's Research park in Eindhoven.

The industry uses this infrastructure to enhance their own methods of handling their own scientific data. And even Unilever which is a big food and chemistry concern in the Netherlands uses it to handle their food informatics data. They use our E-science technology to go through their food information databases to answer questions like: in what ways do which people taste what they eat? Phillips uses it in their medical department where they have made medical implementations and also use remote experimentation technology to enhance their products. **As a result we have close working relationships with several large companies many of which are also working together with SURFnet in university-based private sector partnerships.**

COOK Report: You mention E-science and databases and the optical network and I wonder how you would describe the intentions and resources that the commercial companies and universities bring with them when they sit down at the table to talk about working together. Do they just look for the ability of their staff scientists to have general day to day intellectual relationships with their university counterparts? Or does it get more specific and do they say we're looking for a new kind of algorithm to do some specific task what talents can you bring to bear it and perhaps contract-based research to help us achieve this end?

De Laat: They develop work packages and within those work packages there are tasks for the different partners. As a result they can build a very formal project.

But what industry really finds interesting is that we bring to them a basic infrastructure based on a grid computer architecture enabled by high-speed optical networks and containing a layer of very useful middleware. In the middleware layer you have software controlling the networks but the same of course holds true for the grid middleware layer that includes Globus for applications scheduling and distribution.

One of the things that industry is typically very interested in is the data explosion problem. This is a focal point that by means of our grid infrastructure and parallel processing software we can help them successfully tackle. They need to share data with each other and to do data mining of the data they create. You need Web 2.0 technology and metadata capabilities to describe the raw data and make it more accessible. This also needs ontologies and, as a result, one of the group's at VL-e is working on needed ontologies to describe tastes in a food database from Unilever. If you want to do data mining and cross correlation between groups of people with tastes, this infrastructure is really necessary to obtain that goal.

COOK Report: I think I am hearing from you that the truth of the matter goes far beyond the assumption that, with the big pipe connecting corporations with researchers and universities with researchers, the scientists on both ends would figure out some way by which they could collaborate with each other.

De Laat: We do things differently. Here is a slide [see bottom of next page] that describes the infrastructure that we have built to enable serious ongoing E-science research cooperation between universities and enterprises. This is the slide from the slide deck called Beyond Hybrid Networking http://staff.science.uva.nl/~delaat/talks/cdl-2008-11-14-F.

pdf and I'm showing the fourth from the last slide that is about e-science.

At the foundation layer you have a grid substrate which is the physical network. Above that you have the middleware high-performance distributed computing web and grid.

Neggers: Supercomputers, telescopes, microscopes and whatever else you need as the basic tools required for your research are considered to be in the chocolate covered box called high-performance computing and storage. All these resources in the chocolate brown bar are connected to the network to be used by the users. There is software needed as a layer between the network resources and the users to enable the use of the entire package. The users are in different disciplines represented by these smaller vertical bars and what is becoming much more widely acknowledged now is that the most attractive areas of research are not within these disciplinary silos but rather within the interdisciplinary spaces between them.

De Laat: It is the combination of these disciplines that gives new science.

Neggers: And it is in these little dark areas between the disciplinary silos that the new inventions take place.

COOK Report: This is very interesting because the approach that I wrote about last summer was much more narrowly focused in the belief that a company would identify a specific problem to solve. It then needed to go to a specific researcher and develop a narrowly focused contract applicable to the task and that, somehow in advance, everyone on both sides of the fence would know what was needed.

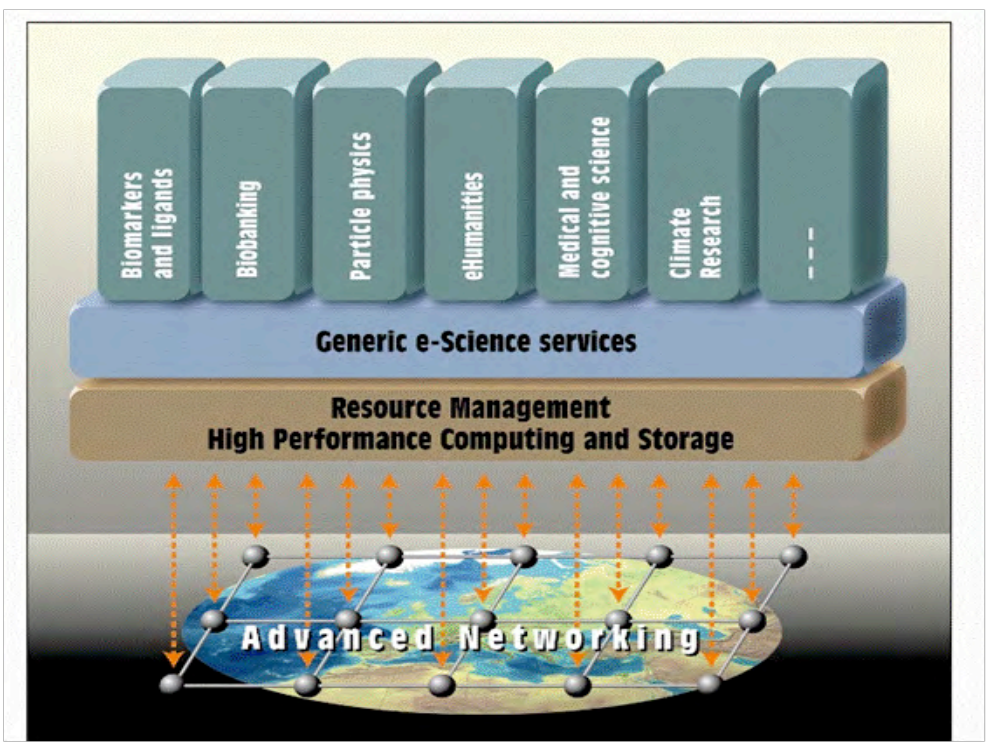

Virtual laboratory for e-Science, *Bob Hertzberger*, *Henri Bal* et. al.

The Optical Network and its Platform for Collaboration Become Basic Infrastructure

Neggers: But more important for the Netherlands, and switching back the focus to us, is that the GigaPort project is now declared a large success by international reviewer's and by our own users. Our government has been informed by the reviewer's and especially by the Dutch Roadmap Committee for Large Scale Research Facilities, which advises the Dutch Government on our participation in the European wide ESFRI projects that this network is needed in the future regardless.

They were told that an advanced network is necessary no matter what you wish to do and they have been given the advice now that SURFnet should no longer be funded out of project funding but rather

basic infrastructural funding is needed to keep SURFnet in place.

It seems very significant that the government has recognized this and the Ministry of Science and Education has now prepared structural funding on an ongoing basis for SURFnet's continued role in scientific and technology innovation in the Netherlands. **But even more significant is that they are saying that SURFnet is not enough and that we also need software to glue this together and the ICT components that have enabled SURFnet and the VL-e project to succeed need also to be funded in the future on a structural basis and no longer on a temporary project basis, where they allways would have to compete with normal research projects.**

Therefore the Minister has asked ICTRegie to make a 10 year forecast for the necessary integration of the advanced network and the advanced ICT infrastructure that enables it to perform its fundamentally important role supporting the innovation necessary to continued economic growth and development. This means, that in the future researchers no matter whether they are in universities or corporations, will be able to rely on the high-quality ICT infrastructure needed that includes a hybrid network, high performance computing and all the storage you would need as a foundation for their further work.

COOK Report: How much understanding of what you have just described to me exists here in the exhibition hall at the Supercomputing 2008 meeting that is the premier high-performance computing meeting in the world?

Neggers: Some but probably not enough. **The key issue is that governments need to understand this and not just the practitioners at this meeting.**

We consider this ICT infrastructure up to and including the blue layer here not just as basic network infrastructure but rather as a more general infrastructure for the information era. Just as you needed the railroads in the 19th century, we need this now and having it will create new jobs and new growth.

De Laat: Hardware is very important for the chocolate brown layer in the middle, but for the blue layer above that, you need people.

Neggers: **In our view innovation needs the network and hardware resources plus the software and the people that have the knowledge of how all of these layers interact with each other and in addition to this knowledge you still need a major outreach effort to involve the users for which it has been created.**

De Laat: **From the light blue box you need people working with every discipline to make it happen, because the people in the disciplines themselves cannot do that. You can not expect the biology professor to understand how to go out and identify and hire the people necessary to teach him and his colleagues how to use the high performance tools.**

COOK Report: In the United States if you promote technology transfer from research networks to enterprises, and do it without an infrastructure like ICTRegie, you will be dealing with professorial fiefdoms and the kingdoms of University incubators on what might be perceived as a competitive basis.

See the article from the *New York Times* found here: http://www.nytimes.com/2008/09/07/technology/07unbox.html?_r=1&oref=slogin&pagewanted=all Is this a problem in the Netherlands?

De Laat: Not as much I think because we are smaller we have only about 12 universities and a number of technical institutions and we tend not to have universities with vast private endowments. In the Netherlands however we also have developed specific groups that are working with the kinds of computer science that we developed to assist disciplines in applying it to their research priorities.

Neggers: We do have all this in place. We do have a hybrid network connecting all the institutions. We do have a high-performance computing budget that renews supercomputers every four years. We do have a large storage grid. We do have E-science tools for laboratories. All this is already in place.

The new thing is that the government now prepares to finance all of this on a structural basis and the only condition they will place on this is that the research community as a whole in the Netherlands will create a national governance structure that will make this a well integrated effort where all the components will be in balance with each other.

COOK Report: What are the disciplines listed in the small vertical cylinders?

Neggers They are genomic markers, biobanking, fundamental research into matter, virtual knowledge studio, cognitive science, environmental disciplines earth sciences oceanography and the final cylinder says "essentially and so on". This is not meant to be an all embracing list of disciplines.

Fortunately there is complete understanding within our government how important this is to maintain a modern economy.

De Laat: The journal that I gave you is a special issue about the application of the Optiportal to this E- science infrastructure.

Larry Smarr, Maxine Brown, Cees de Laat, Editorial: "Special Section: OptIPlanet - The OptIPuter Global Collaboratory", *FGCS*, Vol 25, issue 2, feb 2009, pages 109-113

L. Smarr, T.A. DeFanti, M.D. Brown and C.T.A.M. de Laat, "iGrid 2005: The Global Lambda Integrated Facility", editorial, iGrid2005 special issue, *Future Generation Computer Systems*, volume 22 issue 8, pp. 849-851 (2006).

Tom DeFanti, Cees de Laat, Joe Mambretti, Kees Neggers, Bill St. Arnaud: "TransLight: a global-scale LambdaGrid for e-science", *Communications of the ACM*, Volume 46 , Issue 11 (November 2003), Pages: 34 - 41

COOK Report I'm thinking that Obama science transition team needs to understand all of this. But of course how to do it is another issue.

Get the Optical Infrastructure Right

Neggers: Get well connected and that point becomes a magnet for further connection. From NetherLight, as you can see on the map at the top of page 14, you can go to any place you want to in Europe, the US, and Asia. We are very well connected as you can see in this section of the latest GLIF map from May of this year. There is also now a research dark fiber from St. Petersburg into Finland and it means we will soon have 10 Gb Lambdas from NetherLight to St. Petersburg that will then be connected into Moscow.

The world is a global community now, anyone wants to peer and work with the very brightest people on the planet no matter what country in which they live. The networks are the vehicle that we build for users and for applications to meet worldwide.

COOK Report: While in theory then an American innova-

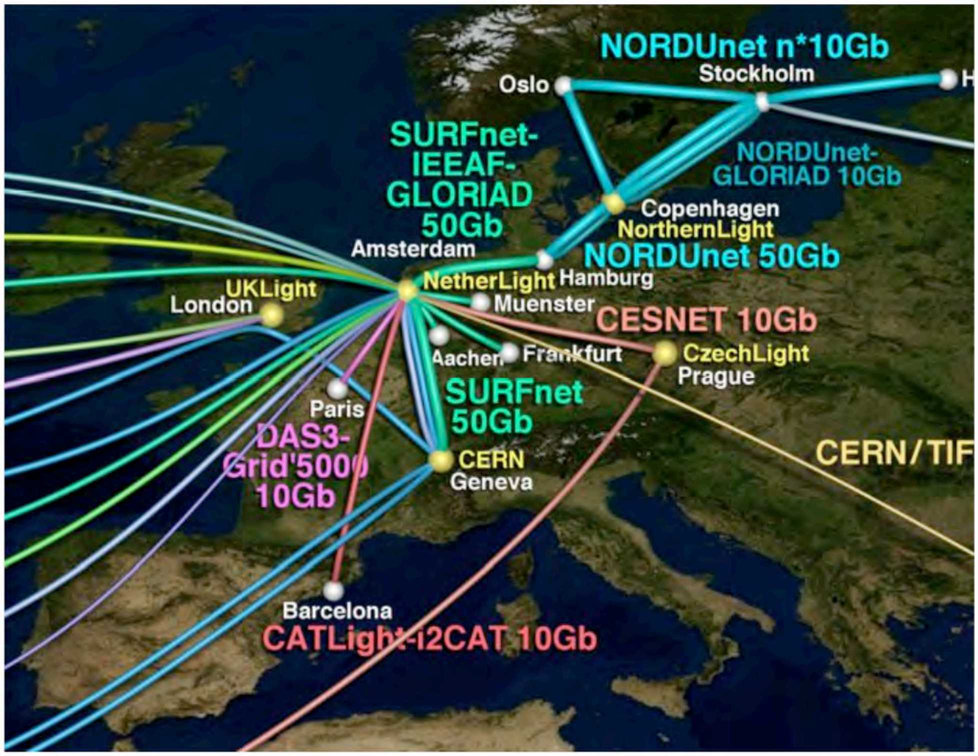

tion network effort focusing on National Lambda Rail could buy a 10 Gb link from MAN LAN in New York City to Netherlight and plug-in to some of the work of your e-science virtual laboratory? What would be necessary to try that?

Kees Neggers at Sc08 in Austin

Neggers: A very simple answer. First of all, NLR is connected today via the many GLIF links connected to NetherLight, including IRNC, IEEAF and GLORIAD links. Secondly, at the moment we are doing a procurement to renew our trans-Atlantic lambdas to New York and to Chicago. NORDUnet and NLR are also taking part in this procurement and will be getting their own links from NetherLight to MAN LAN in New York too. Any American based innovation effort could take a 10 Gb link from New York to Amsterdam for about €100,000 a year. The cost per month of a 10G lambda from Amsterdam to New York is now less than €9000.

COOK Report: But the 10 Gb Lambda for an enterprise would be more than €9000?

Neggers: No, this is a commercial price. It is based on a public procurement. We have had bids from more than 10 different players. So prices between Amsterdam and New York are very competitive. But if you want to go beyond New York the difficulties start. If I want to extend the link from Amsterdam to Chicago, it doubles the price of the link from Amsterdam to New York.

COOK Report**:** And how do you get from NetherLight with a direct connection to AMS-IX?

Neggers: That is very easy, NetherLight is at SARA in the same building as one of the nodes of AMS-IX. And many install lightpaths to AMS-IX now to take advantage of the competitive internet upstream prices in Amsterdam.

The Basis for a Future Internet?

Optical Hybrid Networks and e-Science as Platforms for Innovation and Tech Transfer

Editor's Note: I continued the discussion begun on November 17th with Kees Neggers and Cees de Laat, on the 19th with Cees alone.

COOK Report: Where does your work on optical hybrid networking fit in the context of the E-Science program at the University of Amsterdam and the Dutch e-Science program Virtual Laboratory (VL-e) work we talked about yesterday?

De Laat: Let me try to give you some context. The first slide shows our organization within the University and the topics that my group investigates on:

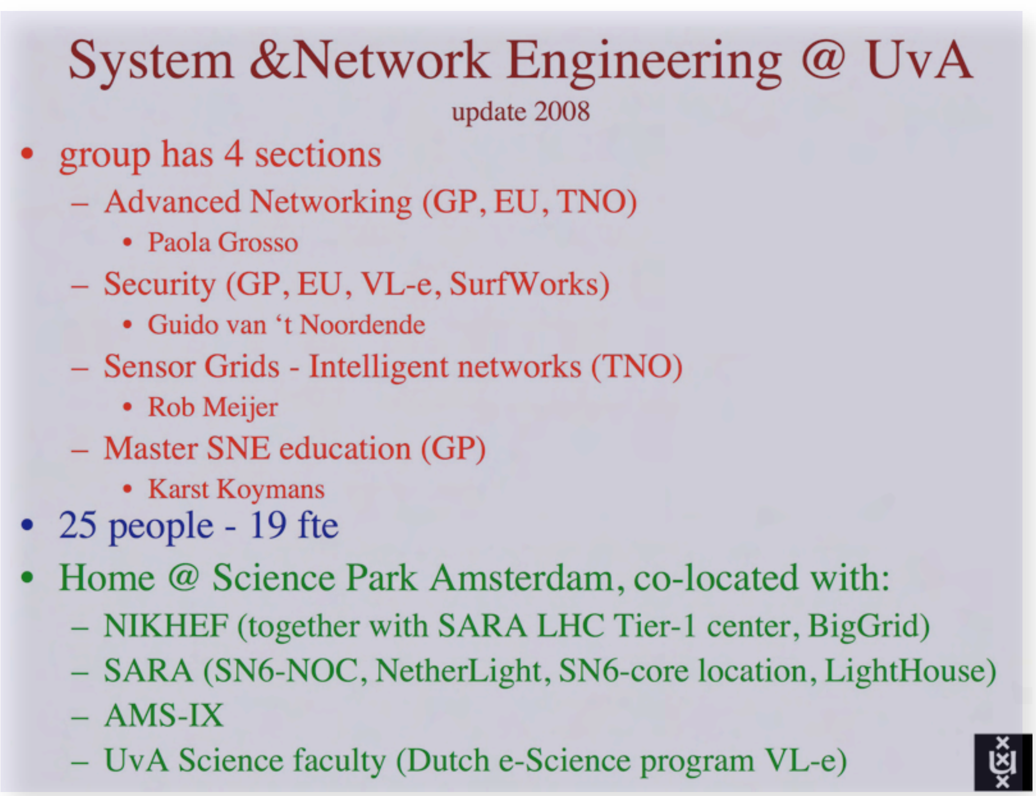

My group has four sections lead by senior scientists. One of the groups is lead by a professor from Industry, Robert Meijer.

Rob started out five years ago with an Internet of "things" – that is sensor networks. He came, like myself, from high energy physics and then started to work in the telecom sector in the research department of KPN. He did a lot of virtualization work on the telecom network; for example you could go to a web service, type in two phone numbers, and have them call each other. Effectively what he was doing was making network components into software objects.

He is approaching this by virtualizing devices and turning them into web services. He encapsulates the network elements into sub routines and then makes them callable from application programs

COOK Report: One of these virtualized devices then is a software interface that you use as a set of instructions plugged into another piece of software? You would have many different software descriptions for different devices such that when you activated the description, it can then go to the device, communicate with the operating system of the device and tell it what to do?

De Laat: Exactly. This is his general approach. He

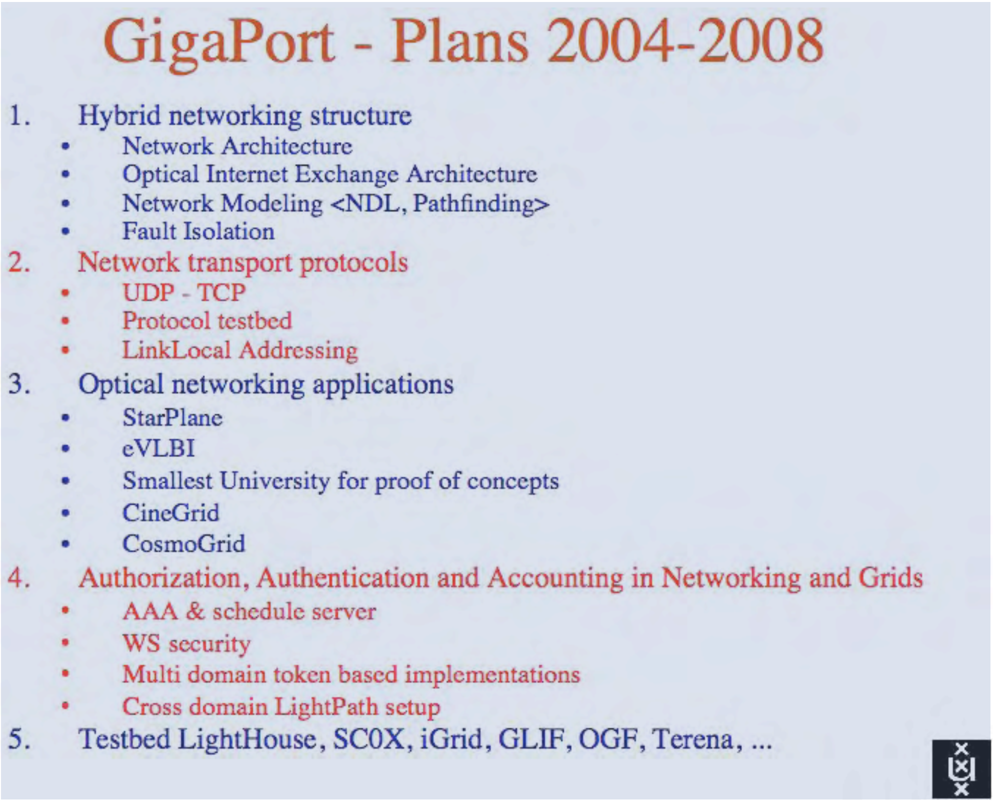

works on it with several (PhD) students. He creates network objects that are basically subroutines such that you can ask them: where are you? how busy are you? and what can you do? You can put a problem in Mathematica that will use these subroutines to seek an optimal solution for the application.

Within a network, for example, you can have subroutines that describe the topology of the network. Then in Mathematica you also have objects that describe the problem space and the formulas needed to solve the problem. Next you put brackets around them and instruct Mathematica to solve and optimize the bracketed formula containing those callouts to the network. Your subroutines undergo a series of transformations in which Mathematica tries to make what actually happens more closely aligned with your desired outcome.

In some sense what I am describing is a kind of peculiar programming language or programming environment where your networks are just subroutines, your data are more subroutines, your solutions to solve the problem are subroutines. You then say optimize and solve and it will then work out the most optimal way to achieve what you want to have done.

What all this means – and this is the most fundamental thing to get one's mind around – is that your network becomes just part of your programming environment.

Normally you have your data and your computing and this stupid network that is an unmodifiable "given" that you have to play along with. **But here your network is just part of your toolset.**

Take the concept of MPI (Message Passing Interface). When you have to calculate a huge "do-loop," then, depending on the calculation, you can divide the work amongst a number of different processors. For example one can let ten processors each do one-tenth of the work; they can work in parallel and be ready in one-tenth of the original time. MPI is a programming library that takes care of distributing the data and instructions among the different processors. We see a need and possibility for a similar framework for networks. This then becomes part of the toolset for creating the necessary network subroutines that you can use to make a kind of automated construction of your desired network usage into a part of your programming environment.

As such the envisioned toolset is a programming environment to optimize the data and communications used for your networked applications across multiple blades of CPUs across the (wide area) network. Just like MPI gives you loop unrolling and parallelization you can also do communications allocation and parallelization so that your application gets a deterministic network to operate on.

You grab a piece of the communications hardware and optimize it for the problem that you have at hand.

COOK Report: What do you mean by unrolling in the network case?

De Laat: Manipulate a network for your purposes. Suppose that you have a huge number of network links. They are sitting there but they are not doing anything for your application. But what would happen if you could manipulate them? If you could realign them? If you could take them and put them at a different co-location point? Say for example, my data center is here, but I want my photonics to go over there. If you have the ability to reconfigure these kinds of resources, you then get the ability to optimize the network tools at your disposal.

Your application does not need to use everything that is there. You can take just those resources that are useful for a particular purpose and twiddle them around and solve your problem doing so ten times as fast in a deterministic way. **At that moment your network becomes part of your programming language and your problem-solving environment**.

How the World Changes When You Give the User his Own Network

Now we are taking some baby steps down this road. And we have developed some forwarding engines and some device engines to manipulate tables into packet inspection and add headers to packets so they can be manipulated, mostly in software, in the control and routing portion of this network.

In our group back home we are creating objects for virtualizing and programming wavelength switches and photonic devices which directly talk to fiber; Micro-Electro-Mechanical Systems (MEMS) devices that can connect fibers so that we have flexibility at the fiber layer. And at the Ethernet layer, we can do similar things manipulating Virtual Local Area Networks (VLANs). We address the Ethernet layer and the packet routing when we need that.

If you can manipulate all these layers and have also vertical and horizontal knowledge in every layer, you can do the magic and you get a perfectly integrated multilayer hybrid network that is optimized for your application.

That is what we are embarking on and these are the first baby steps. My expectation is that, given funding for our work, we are moving forward and we have further research ideas to extend this technology. For example this will extend our topology, policy, and authorization and authentication development.

While your own part of the network functions, the other parts may not. Therefore, while you are optimizing your network, you will need to take into account the different usage policies of the networks, that is to say who can use the network under what conditions. One also needs to address the cost. Our work has several components. A policy component; a path finding program, an optimization component and a cost finding function.

Cook Report: if we ever get to this future world where we either get rid of the phone companies or transform them and we have fiber everywhere, then does it become possible first to have these capabilities available to every fiber connected household in a city like Amsterdam for example? But then can we take your thinking a step further?

You will ultimately have a vast range of possible things that can be done and you will likely be thinking of policies that will authorize appropriate users to appropriate levels of tasks. You don't let just anyone into a nuclear research facility obviously, but into things like weather simulation modules might you authorize almost anyone who was interested and could use the technology in a self trained or self-motivated way? Could such users include all ranges of the public from grade schools to high school graduates to university graduates to retired people to people from technical schools – you name it?

De Laat: I can answer your question in two ways. The resources that I have explained to you certainly represent value and we can play around with them as we have been doing in a fairly small research setting. But if we do not solve the policy problems of authentication and authorization, then we will never successfully roll it out because the people who own the resources with which these tools will enable experimentation will demand some degree of control over who has access to them. To some extent to determine access, you begin to build in telecom technology again, but on a different level. Here it is to enable users, not to cut up bandwidth and drive price by making it scarce.

COOK Report: You have a universe of resources that you want to make available under certain conditions to a universe of users. What I am hearing is that under the right conditions a policy administrator could work toward the goal of making available appropriate resources to appropriate users?

De Laat: Not quite. What you need to give the owner of the resource is the ability to set his own conditions. If he wants to give them away for free for example he can do so. But if you don't solve that problem the resource owning people will be very wary because they will complain that you will also have a bunch of people who will want to do illegal movie copying and that sort of thing.

Now there is another important aspect. **We want to be able to give complete lambdas, in other words complete colors, to applications, therefore, the cost models that support it will depend on the underlying technologies used.** The costs involved must be very scalable all the way up toward very high capacity. And this is why we are especially focusing on switched photonics and avoiding where possible much more expensive routing. In routing the cost factors are different but also these services are different.

Given these considerations, you need to do things on the lowest level of the protocol stack which avails you of the service that you need. If I have five people using a printer I put them on the Ethernet level because otherwise it would be a nightmare for the users to get to the printer. Or should I go to the fiber and say: you need to flip that mirror before you print the document otherwise the printer cannot see your document. You go to the layer that is optimal for the problem at hand. No lower but also no higher.

COOK Report: Consequently, at some point in the future if you have someone who has developed an application that could make use of these kinds of network resources, you would be able to put into the hands of that developer some tools by which he could, if he so chose, contribute what he had developed to the kind of "universal" network administration system that would enter it - like a book - into the library of this global network where it could become accessible to be designated classes of users that the developer desired?

De Laat: Yes, or for example if you want to do 4 k streaming to Lecture Hall B because you are about to give a presentation there on fluid dynamics, you tell the system to make the server available with a lightpath to the display screen in the lecture hall where you will teach.

COOK Report: The number of possibilities are unlimited.

De Laat: For sure. But what we need to do and what I was alluding to is to make this process multilayer so that you can use the layer appropriate to your task. And if you need to work with five printers that is one layer; if you want to stream video to a lecture hall that is another layer and if you wanted to use a browser, then I need to be able to direct your resource request to the normal TCP/IP routed layer of the Internet.

What you are looking at is the user interface but what is really important is what is below the interface.

COOK Report: Given then this world in which network tools are virtualized on many different levels. Can we go then to the rest of your hybrid networking presentation and the extent into which it fits into the E-science tech transfer conversation we had with Kees Neggers two days ago?

I'm looking to continue to grasp this in more detail from an understanding of its potential impacts on enterprise business and education and of course I am also wondering what the business model will be because obviously nothing comes for free.

De Laat: In the diagram above, the gray lattice mesh is the substrate that connects everything together. On the top of the substrate in the brown bar you see a High performance Computing and Storage and Resource Management. This is where all your computing resources are found: data centers, supercomputers, clusters, visualization, measurement instruments and so on.

On top of that in the light blue bar is a middleware layer that virtualizes all of these resources and with grid technology allows people to build services and applications from the virtualized resources. This is pretty much what is already happening in industry as well; here it is designed to enhance science. And on top of everything we have a number of disciplines like biomarkers, bio-banking, particle physics, eHumanities, medical and cognitive science, climate research and many other disciplines.

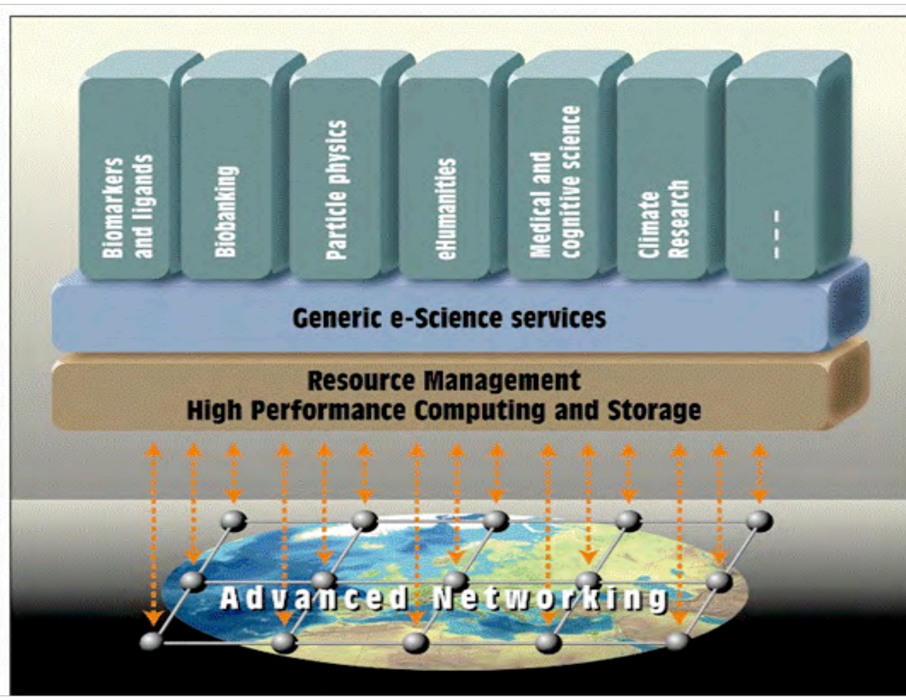

Virtual laboratory for e-Science, *Bob Hertzberger, Henri Bal* et. al.

Figure 1: The E-Science Tech Transfer World of the Dutch National Innovation Network

There is also software in here that is tailored toward each domain of science but still generic to that domain such that different groups in the discipline can share it. And then on top of that you have the different domain sciences.

COOK Report: And when did this start?

De Laat: This method of a systems approach to do science started about 5 years ago with the beginning of the GigaPort and Virtual lab projects. The chart in Figure 1 is really a model containing several projects that are interconnected and shown in relationship to one another. They were all funded around 2004.

The GigaPort project designed the network that is the gray lattice mesh in the picture. The brown bar is primarily the hardware infrastructure and its direct control software (and virtualization) whereas the blue gray bar represents primarily software. To do modern system level science you need a combination of resources, instruments, data sources and to enable a group of different scientists to work together to produce a desired result. The blue-gray middleware allows you to do your science in a workflow environment. In other words they need to gather data, need to process it here; need to store it there; need to transform it over there and analyze it over elsewhere. At the end of this process you will get graphic output, interpret it, draw your conclusions and write a paper. Workflow middleware helps to automate this basic process of science.

Treating the GigaPort as Tech Transfer Infrastructure

COOK Report: Would you tell me a bit more about what happened to cause the government to decide to agree to support this as infrastructure?

De Laat: In order to have successful science, you need a set of tools; supercomputers, a data processing grid, a network, but also common middleware. The common of middleware ensures that not every scientist will have to dig in and reinvent the wheels so to speak. This virtual laboratory e-science middleware is regarded as a kind of tool set that you need to provide to the scientists so that they can do their work. In that sense you should regard the generic e-Science services as an infrastructure for the scientists. The government realizes this.

COOK Report: And has agreed to support it on an ongoing basis?

De Laat: Currently, the government supports Virtual Laboratory on a project basis. For the last 10 or 12 years the government has been funding these kinds of projects. Actually the money to do this comes from a different Netherlands infrastructure funding, named ICES/KIS and FES. The Netherlands has its own natural gas wells. A part of the revenues of this natural gas is put in the mentioned funding schemes to support infrastructure.

Infrastructure, as understood from the point of view of the history of the natural gas reserve funds, has been highways, railroads, waterways, harbors, and dikes – thus large scale infrastructure. As a part of the basis for the last round of funding for GigaPort in an article to a national newspaper Jaap van Till, Felipe Rodriquez en Erik Huizer (*) suggested in august 1997 that this money could also be used for ICT infrastructure. They suggested that for the cost of one highway crossing the Netherlands could have the most advanced research network infrastructure. Then it was decided to start funding digital highways. (*) http://www.nrc.nl/W2/Nieuws/1997/08/21/Med/06.html (in Dutch)

COOK Report: Makes sense. I wish we would do something similar in the USA. But would you go on to talk me through hybrid networking.

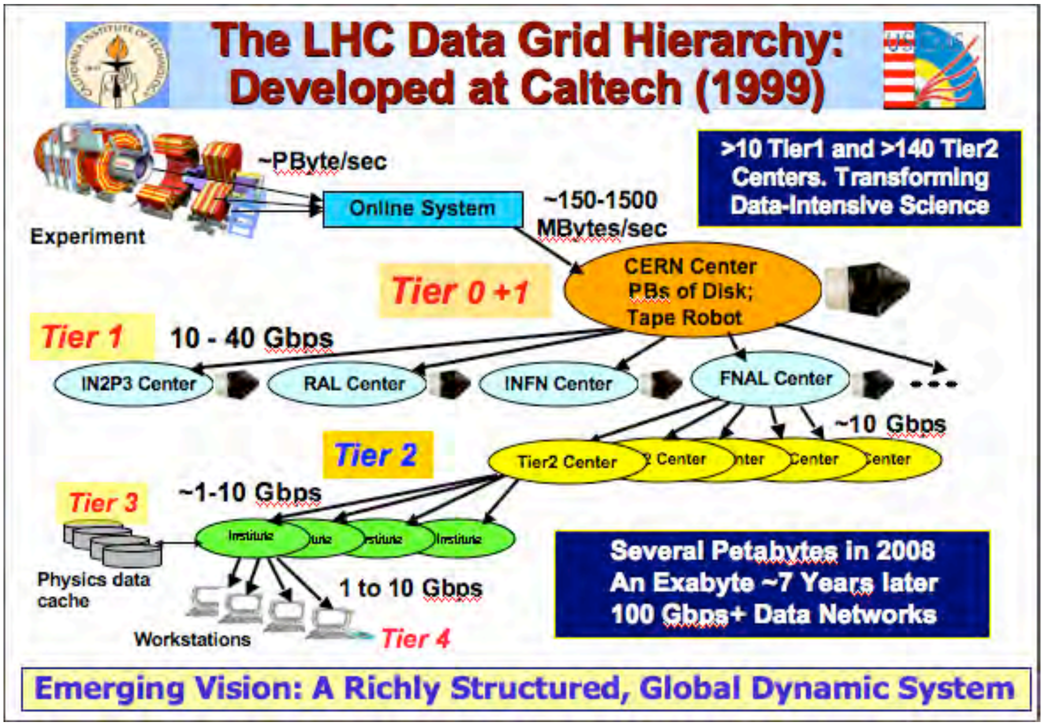

De Laat: Let me do it then from a slightly different point of view with a November 29 2007 talk at the University of Essex called Lambda Grid developments History Present Future - found at http://www.science.uva.nl/~delaat/talks.html

Lambda Grid Developments and the Need for Hybrid Networking

The first need is illustrated by the LHC Network. You have a scientific instrument that generates data in the amount of Gigabytes per second creating thousands of terabytes of data per year that must be distributed to tier-1 centers for processing. These include NIKHEF in Amsterdam, Fermi Lab in the United States, German, English, Italian, and Chinese centers and elsewhere. The tier-1 centers either process data themselves and/or send the data to tier-2 locations and local universities that will process and interpret the data and use the results to do science. Results are then send back to the tier-1's for archival and access by scientists purposes. [**Editor**: The slide shown is from Harvey Newman's September 2008 "deck."]

Consider the next slide shown immediately below: -- Data Analysis at the University of Tokyo. Talk to Professor Kei Hiraki, who has done a number of projects that couple scientific instruments to huge databases across Japan and use computational power to do data mining on these databases using very high speed networks.

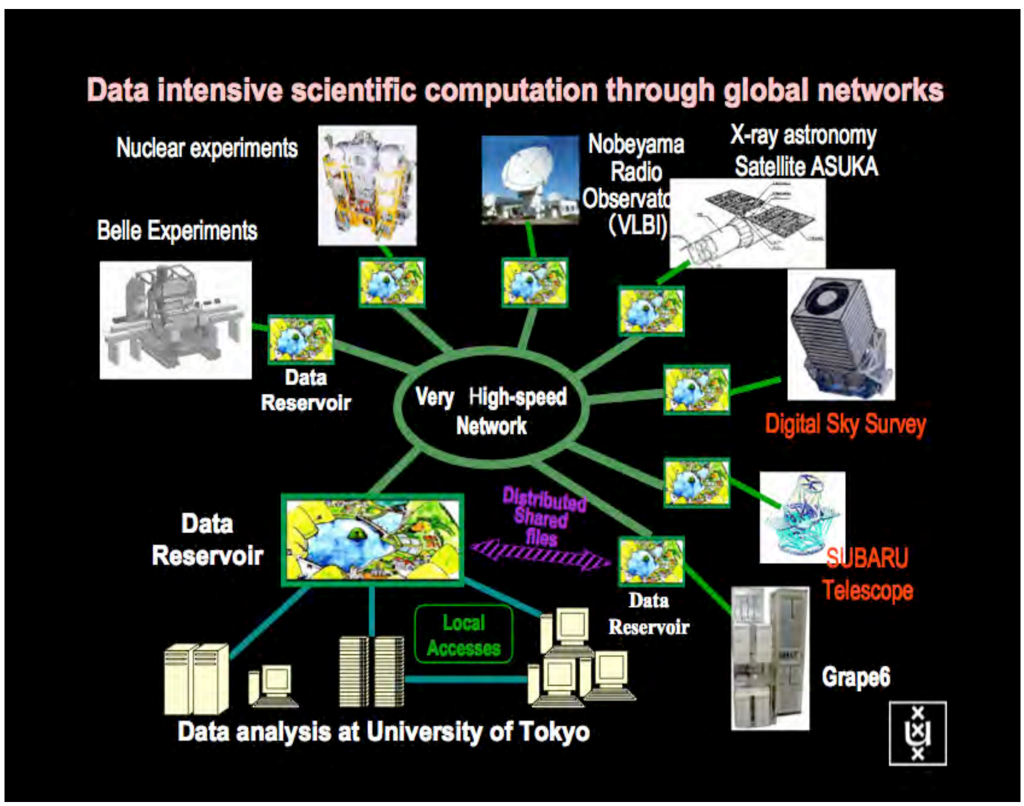

WiVR: a Window into Virtual Reality

The nicknamed "Dead Cat demo" (done at SC04 and at iGrid2005) allows the viewer to explore the internal structure of a physical dead cat (present behind the tablet) using virtual technology. The object was scanned in a CT-scanner in collaboration with the Academic Medical Center at the University of Amsterdam. The data was put in the supercomputer to enable rendering of various slices. This system was put online such that when a request with coordinates of a plane was received, the system would render a picture of the requested slice.

On the demo site using a tracking system the tablet and the viewers position and orientation (line of sight) with respect to the object were continuously measured. The position information was sent to the Amsterdam supercomputer where the new slices are calculated and returned to San Diego to the tablet display (only a 0.5-second delay!). As the person rotates the panel he looks through a virtual window into the object the representation of which is stored in the computer. The panel attached to a high-speed network and correlated with the eyes of the person holding it calculates the correct rendering of the object being observed as the position of the panel changes and transmits that rendering in real time back to the observer holding the panel.

The motivation for this project is to test graphical systems and network performance in the context of a medical application attempting to render images in real time. From http://www.calit2.net/events/igrid2005/?p=45
see also
http://www.science.uva.nl/~mscarpa/wivr/

This has significant implications for surgery and for other forms of telemedicine. A scholarly write up has been published in Future Generation Computer Systems volume 22 2006 pages 896-900 under the title Highly Interactive Distributed Visualization.

COOK Report: This can then become a very useful medical diagnostic device?

De Laat: absolutely. You hand out cheap devices for the diagnostic scans while you locate expensive computer, data and rendering infrastructure located elsewhere connected by predictable networks to do the calculations and image processing.

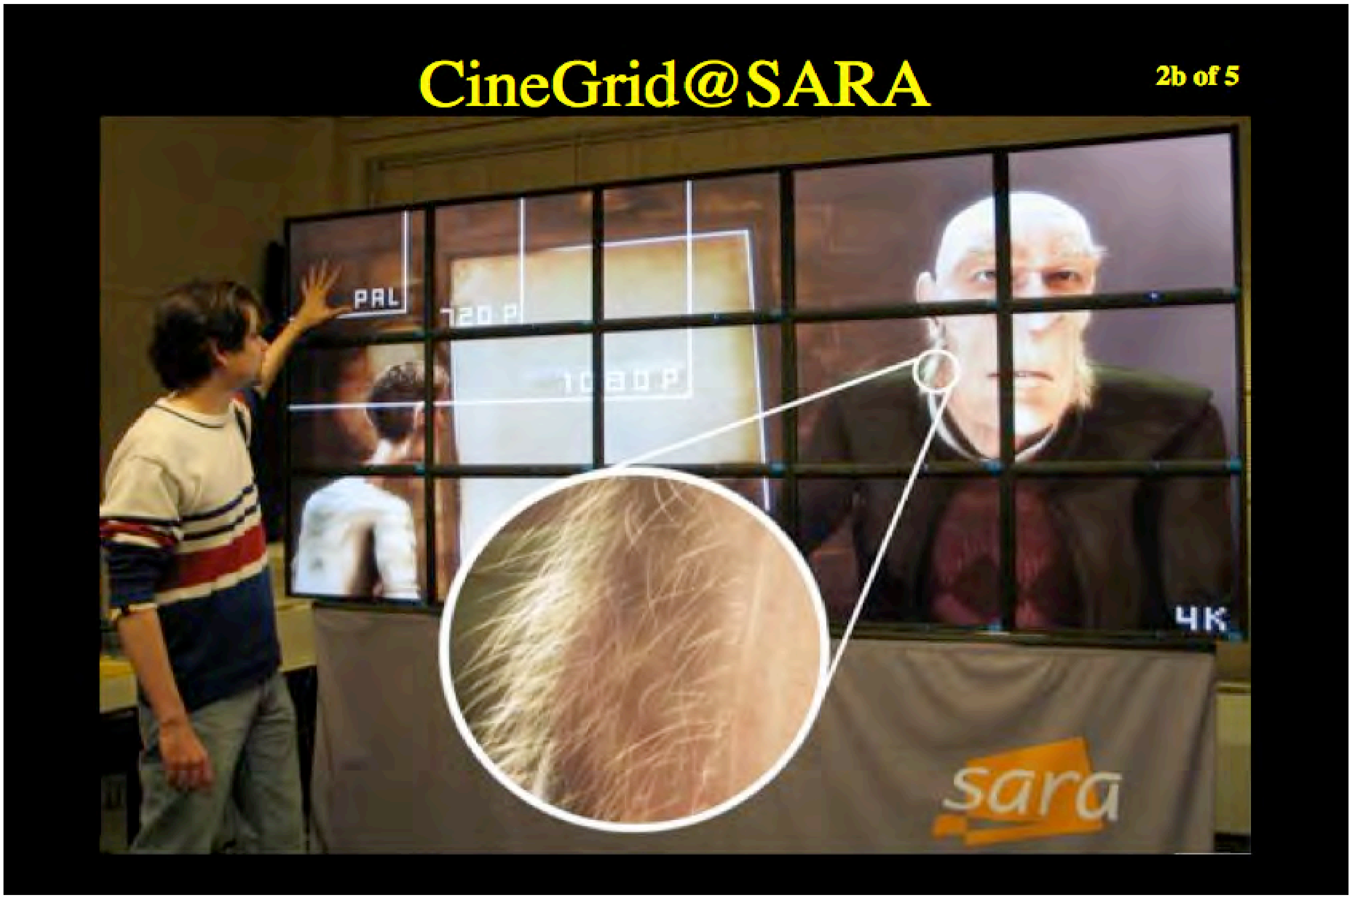

The slide above shows the visualization of cinematic SHD content (3840 by 2160 pixels) on a tiled display setup at the supercomputing center in Amsterdam.

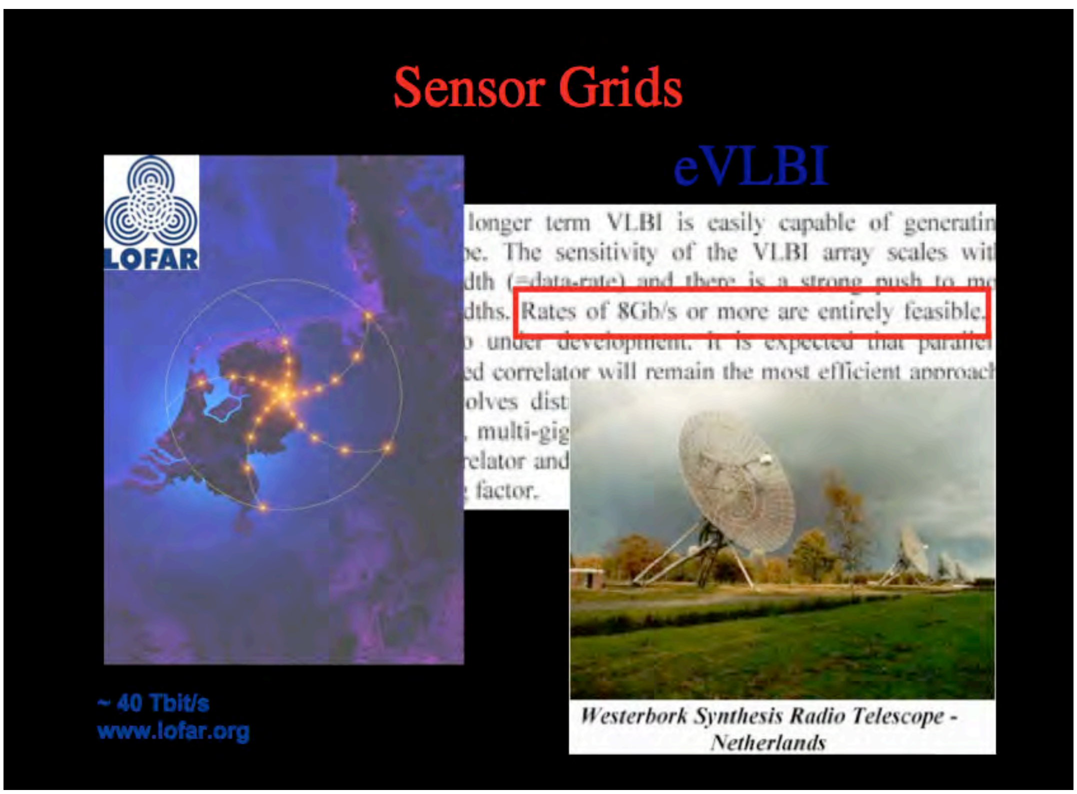

LOFAR: the Very Long Baseline Interferometer radiotelescope array.

The CineGrid Project

CineGrid"s Mission: To build an interdisciplinary community that is focused on the research, development, and demonstration of networked collaborative tools to enable the production, use and exchange of very-high-quality digital media over photonic networks. For more information see http://www.cinegrid.org/

LOFAR: the Very Long

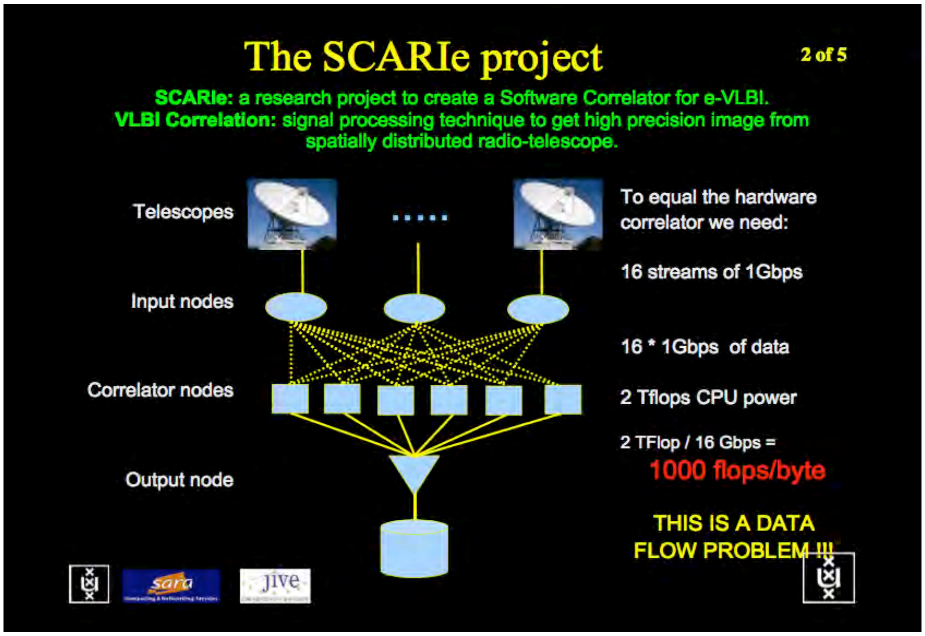

Baseline Interferometer radiotelescope array is shown in the slide at bottom of page 23.

In **the SCARIe project** we use lightpaths and photonic networks to collect all the data from the radio telescopes to put it into a distributed software correlator. The result is as if one is observing the sky with a very large telescope with extraordinary angular resolution.

What These Large Scale Projects Teach Us About User Populations

And below is the ABC user classification slide. It shows how strongly asymmetric the amount of data used by a diverse user population is where the vast majority of people -- your lightweight A class users – need full Internet routing but to do mainly browsing and e-mail and do not use a large amount of bandwidth while in the B class the enterprise users need grid applications, multicast, IP streaming, virtual organizations, mostly composed of LANs. They need VPN services and full Internet uplink routing. Then at the far end of the bandwidth are the most hungry namely the C class of users. These are the users of e-science applications where large data flows exist between a very limited numbers of nodes (e.g. from telescopes to correlator, or from CERN to tier-1 centers). Those flows consume the entire capacity of a lightwave on a fiber and, when let loose over the normal routed Internet would disrupt the operation.

The bandwidth for the class A users tend to be DSL speeds, class B ranging upward to Gigabit Ethernet and the bandwidth needed by the C group is Gigabit Ethernet and above - including multiples of that 10 Gb and 100 Gb. What we need to realize from this is that one size does not fit all. We need to work with a combination of solutions.

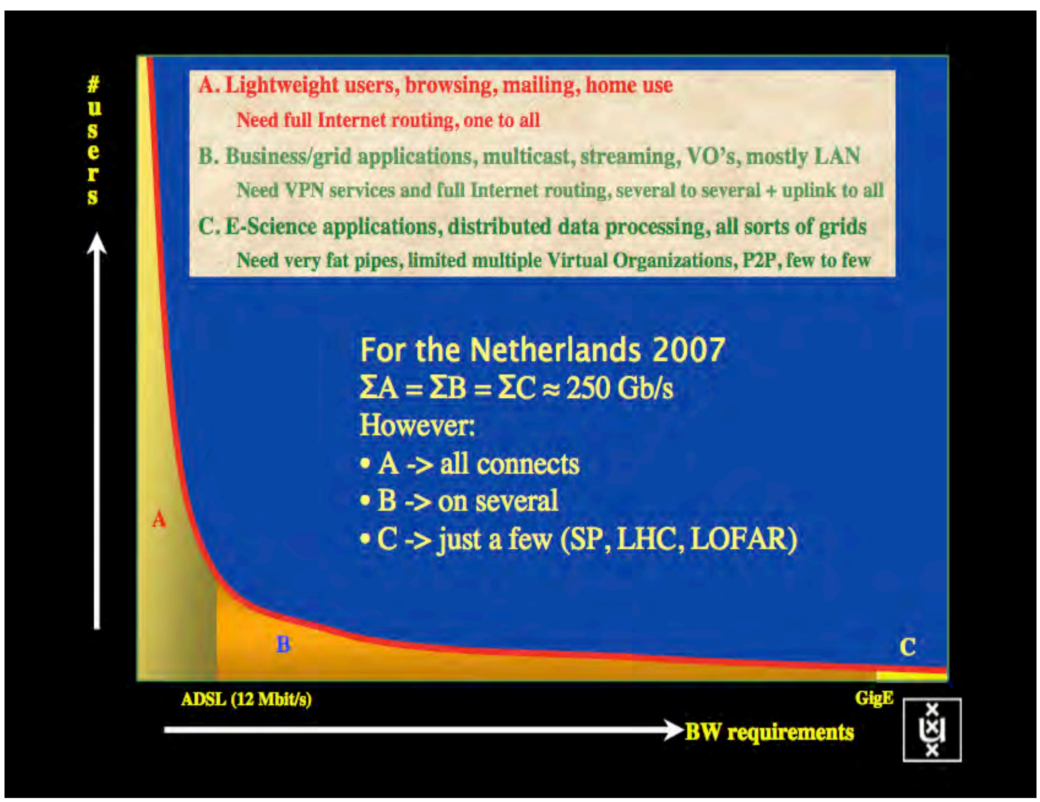

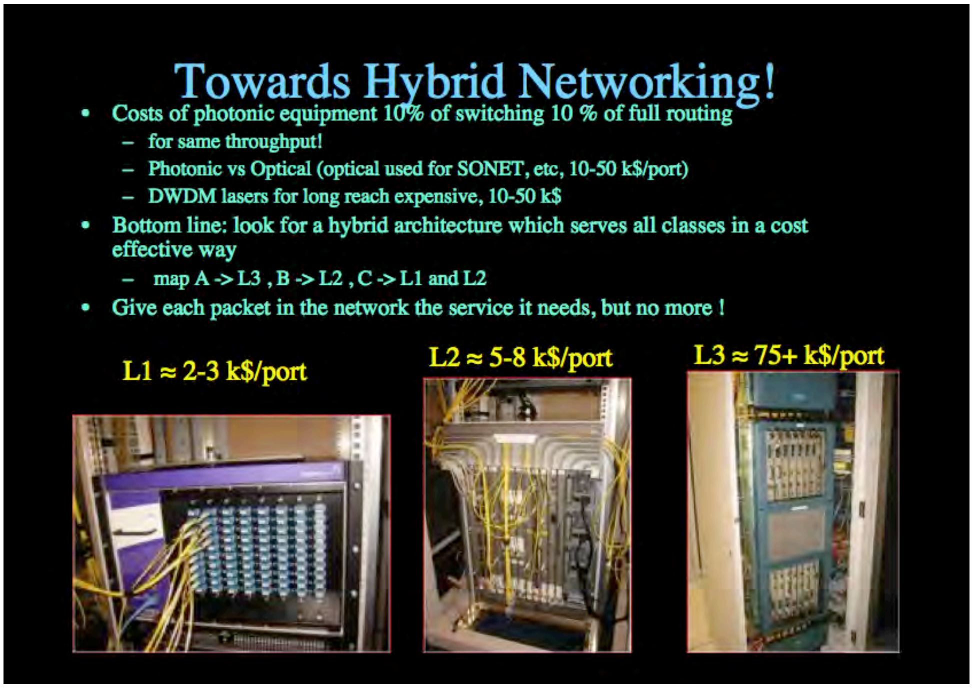

COOK Report: in other words the evolving Internet needs to be a hybridized system that can satisfy all three classes of users enabling transitions from one class to another where necessary?

De Laat: Yes, this was one of the mayor findings in the GigaPort project and this is presented in the slide above "Towards Hybrid Networking."

The cost of photonic equipment has come way down to where it is about 10% of the cost of switching which in turn is about 10% of the cost of full routing. Obviously the services are also different. Routers "know" the topology of the entire Internet and can route per packet, switches using layer 2 information are only aware of the LAN, and the fiber switches only forward entire colors or all light in a fiber toward the outgoing fibers. Using Photonics also takes less energy. But of course if you route you can go specifically where you wish to within the global Internet. The goal is to design your network in a way such that each packet is given the minimal service that it needs to do reach its destination and no more.

For the whole slate of applications on the Internet one will need a combination of the technologies described above. Given the amount of bandwidth that it enables, photonic technology is getting very cheap.

If we begin to switch light waves, we avoid latency and congestion and find that gig Ethernet and above is usually an adequate carrier protocol and that applications can have control over network behaviour that would be very difficult in a routed network.

COOK Report: so what you are saying is that if you are going to invest money in networks you must look at your investment in a very holistic all-inclusive way. And that the appropriate use of your resources would be to match the delivering technologies to the needs of the three classes of users leaving

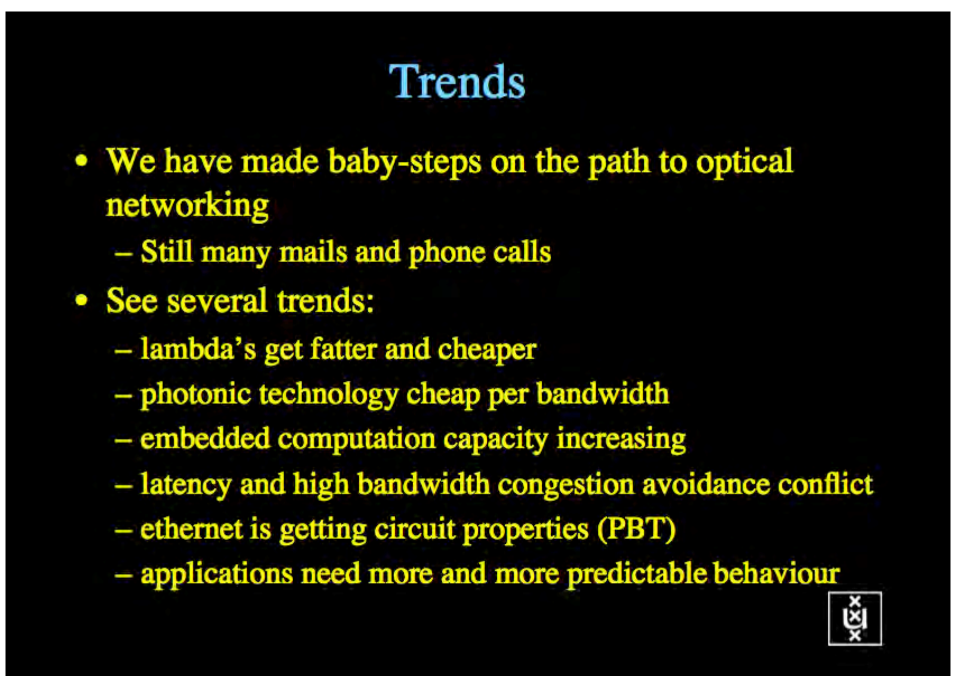

hybrid networks as the logical outcome of this matching of economic investment with technology?

De Laat: Yes just as in computation where one needs to balance grid and supercomputing technology. That is an understanding that we still need to spread. In short there is not one solution that fits everything. You need a combination of solutions for your network transport as well.

How Low Can You Go?

To answer this problem you must ask how low in your stack of network layers can you go? You want to go as low in the service layer stack as possible and go as far as possible at the lowest layer before you need to bounce up higher in the stack in order to do what is necessary to get your packets or photons as the case may be to their final destination.

COOK Report: Then what we started out talking about is **how to build a networking environment that combines this use of equipment, technology, and energy in such a way to give the user applications that meet his demands at the minimal necessary cost?**

De Laat: Yes. But to do this one must expose the network to the users so that they understand what is happening and can see the benefits of not always working at the highest and most expensive layer -- namely the routed layer. One needs to make sure that your users understand the benefits of operating in different layers of their optical network. The users need to grasp that if they endow their applications with the ability to intelligently traverse the layers they can open up all manner of increased possibilities to better performance at less cost. To do this you effectively need a kind of "wall connector" to your optical network represented by the cream -yellow box on the left of the slide above in such a way that the applications can enter the network at the layer on which they can operate most effectively. What I am talking about here is effectively a control or service plane for the network.

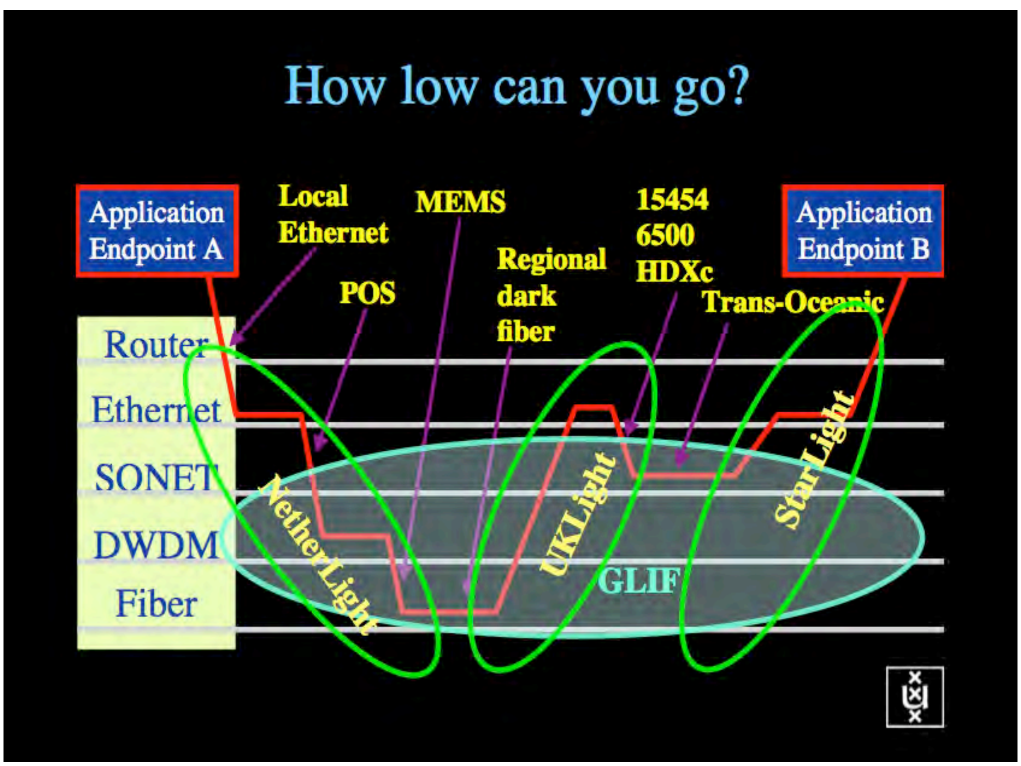

Your Network Application Must Use the Appropriate Transport Layer

The Global Fiber Infrastructure

The chart at the top of the next page represents the GLIF or Global Lambda Integrated Facility. This is the global infrastructure of fiber that we are using for developing the hybrid network tools I am explaining to you.

Also in the next two slides I present the situation around

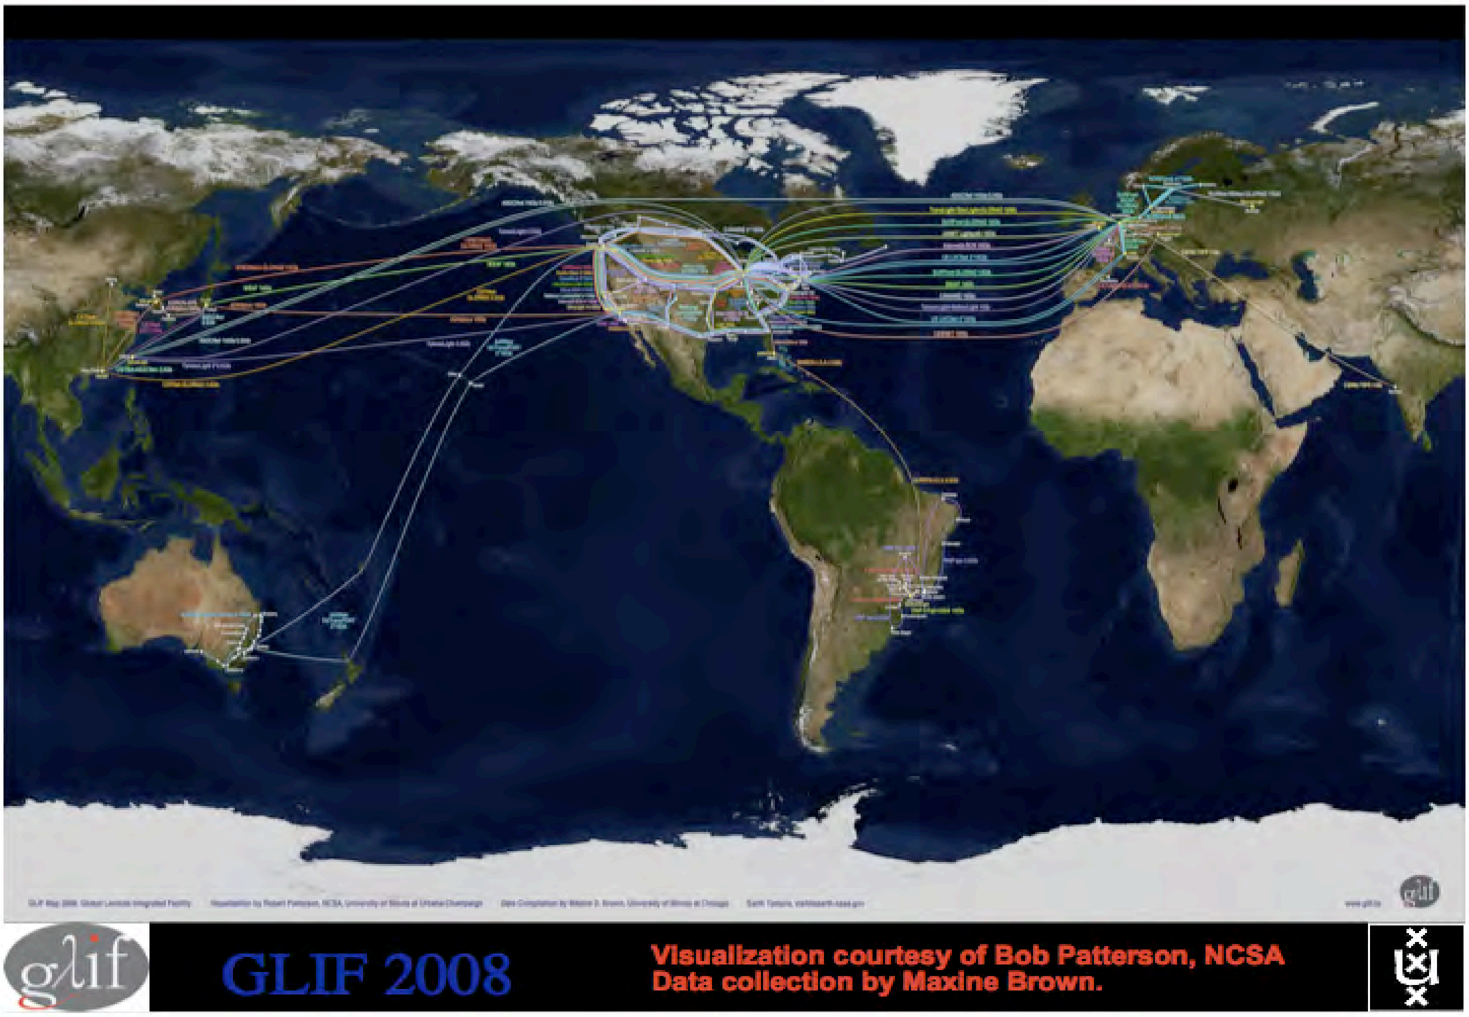

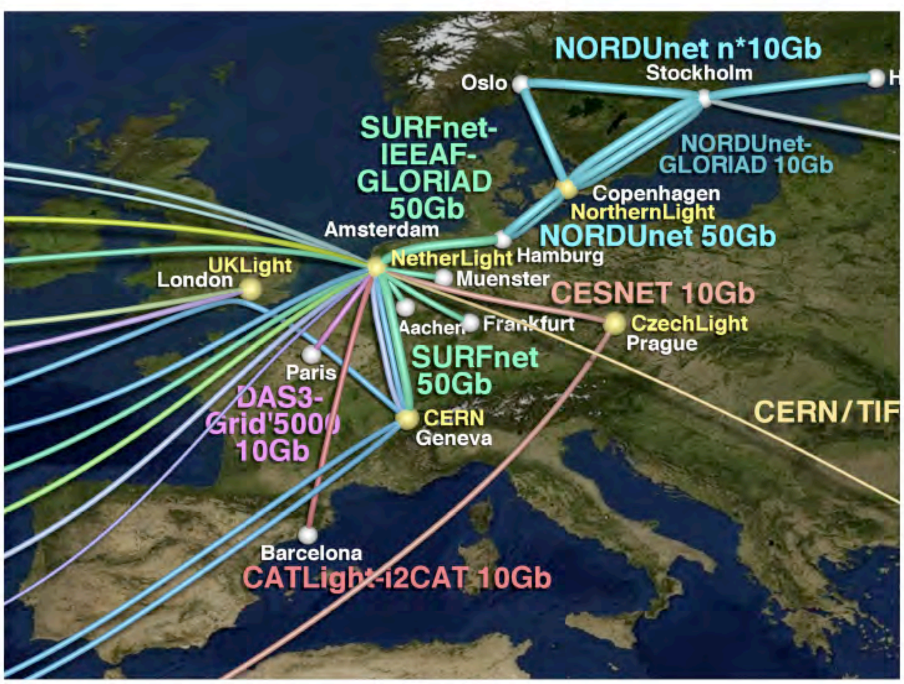

the Netherlands, NetherLight and Amsterdam. First and to the left you can see the connectivity and, in the second graph at the top of the next page and overlaid on the connectivity map, the activities that flourish around NetherLight in Amsterdam.

On page 28 you will see the map of dark fiber in the Netherlands. SURFnet has some 8000 km of fiber - an infrastructure that is larger in length than the railway system of the nation.

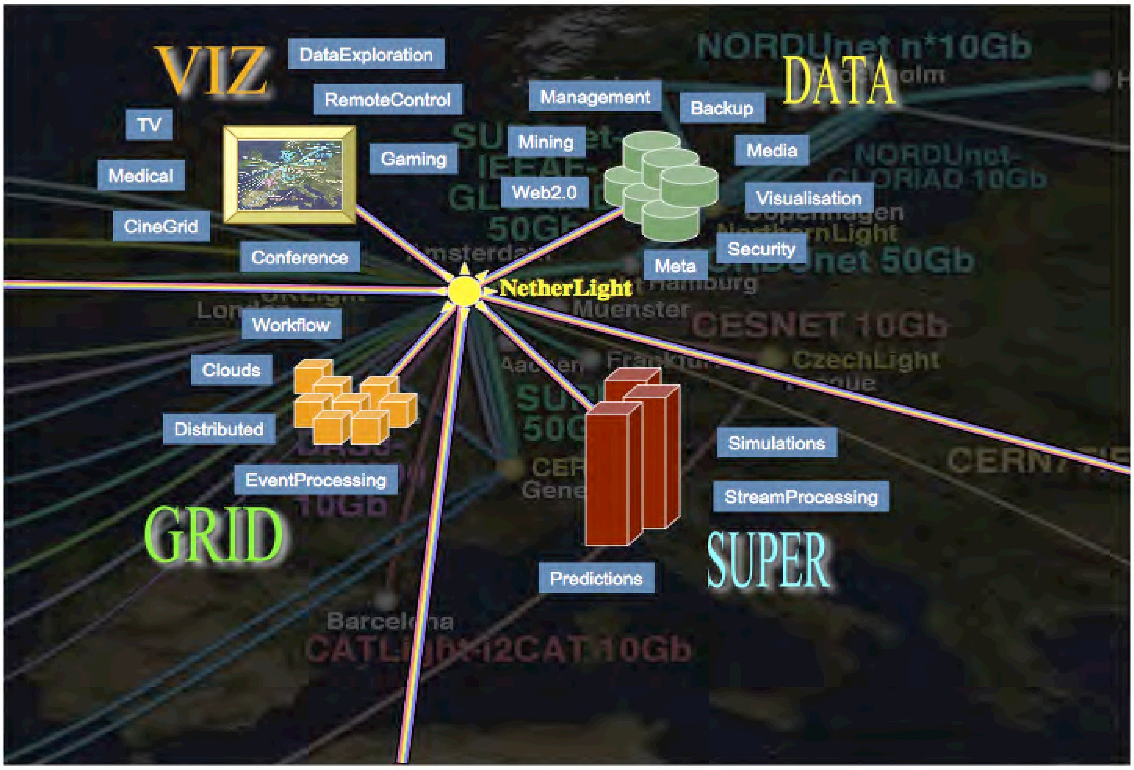

The "map" above shows the four major research areas made possible by the Dutch infrastructure: Visualization, Data Storage and manipulation, grid and cloud computing and supercomputing. The map below shows the SURFnet fiber infrastructure.

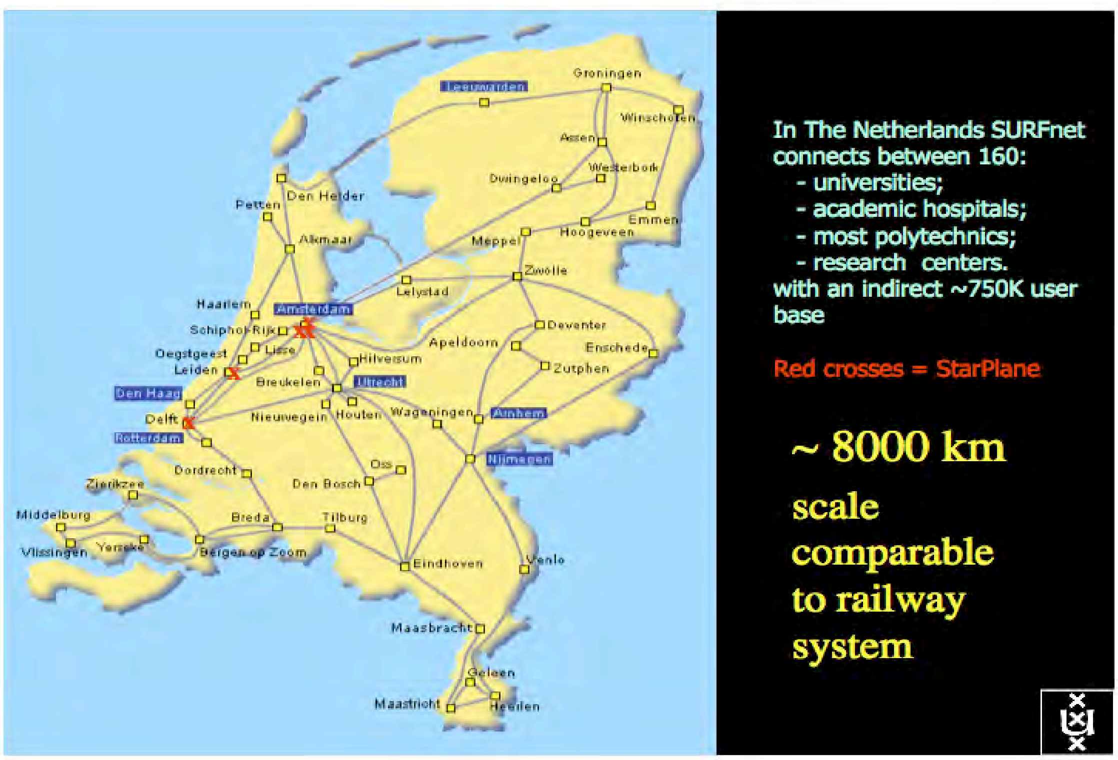

COOK Report: of the 160 connections of SURFnet roughly how many are enterprise research centers?

De Laat: perhaps about four or five

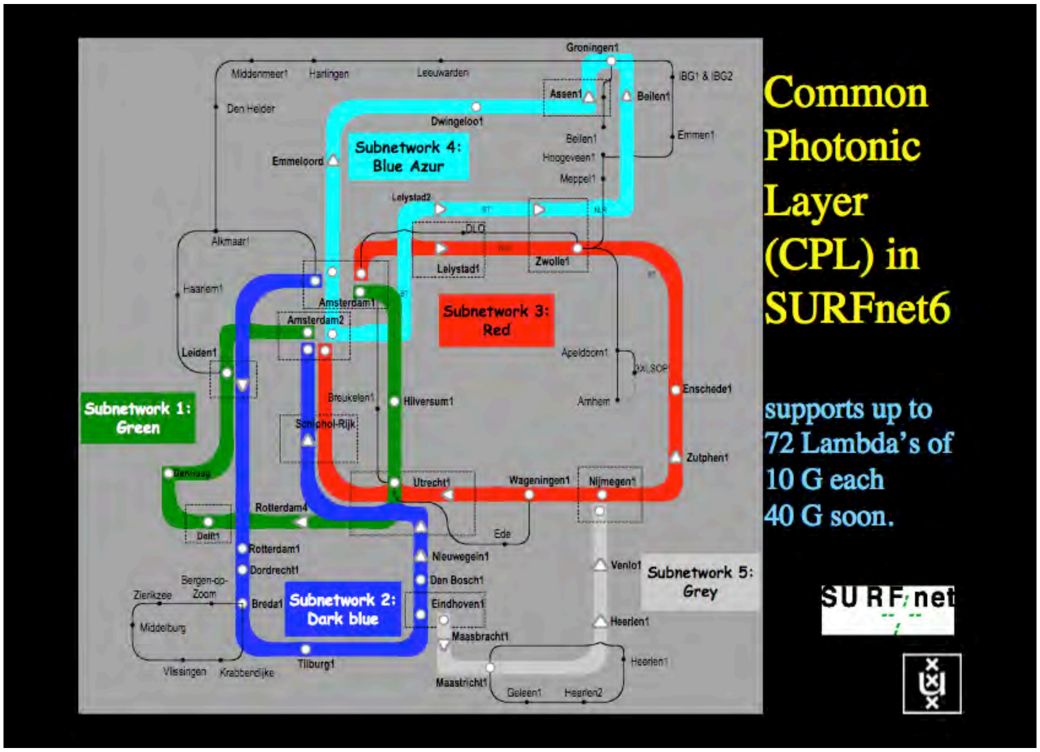

The Significance of SURFnet6 as a Photonic Network

This is the subnetwork map of SURFnet's fiber rings where they can do 10, 40 or 100 Gb per lambda and transport up to 72 lambdas per fiber.

Together with SURFnet we have implemented the StarPlane project with major contributions from Nortel. Its purpose is to allow part of the photonic network infrastructure of SURFnet6 to be manipulated by Grid applications to optimize the performance of specific e-Science applications.

StarPlane uses the physical infrastructure provided by SURFnet6 and the distributed supercomputer DAS-3. Hybrid optical networks such as SURFnet6 allow network users to partition the network resources and to create multiple overlay networks, each with a different logical topology. The novelty of this project is that it does give this flexibility directly to the applications by allowing them to choose the logical interconnection topology in real time, ultimately with subsecond switching times. The novelty of StarPlane is that the change of topology is implemented with photonics; namely Wavelength Selective Switches. So StarPlane rearranges colors in the wide area dark fiber infrastructure. http://www.starplane.org/

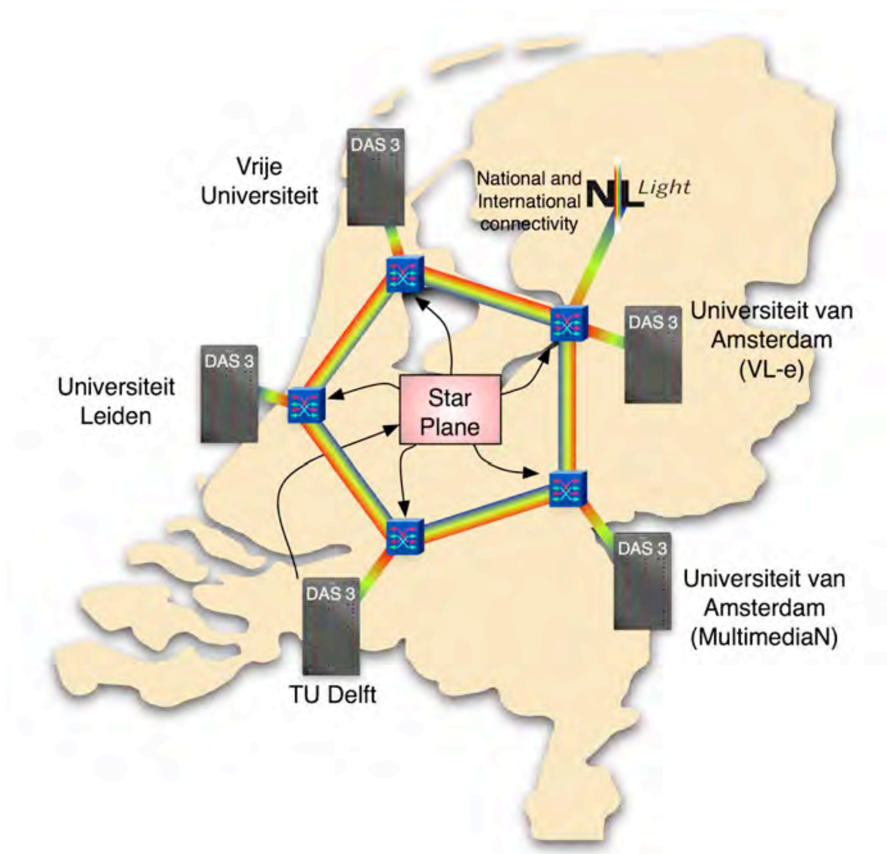

We do this experiment with five distinct clusters at 4 different Universities - hence

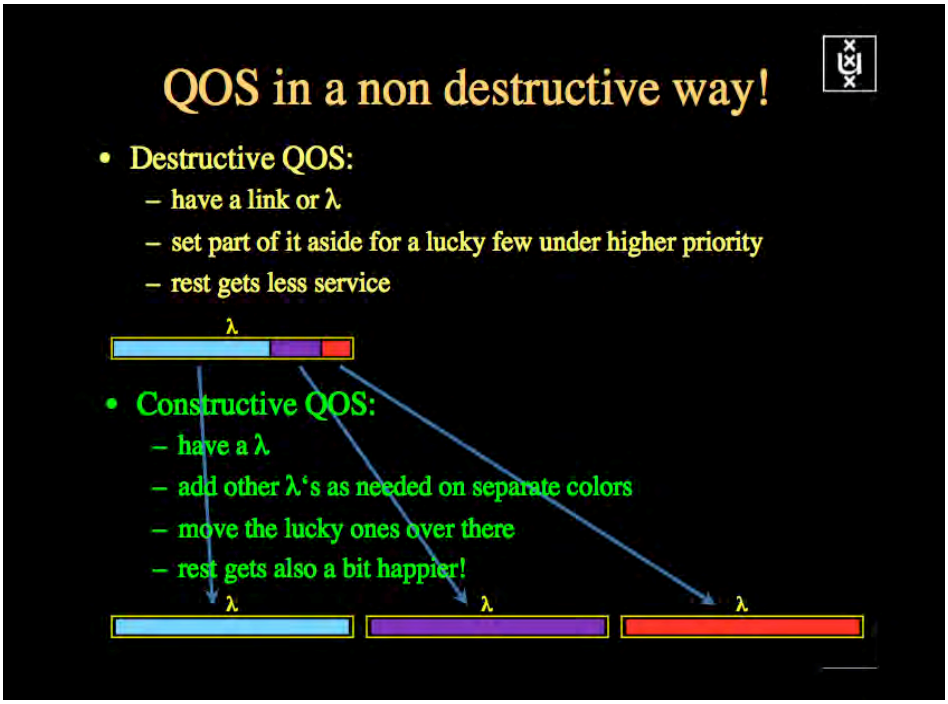

the logical view of a Pentagon shaped diagram. The deployed application signals to the NOC what kind of connectivity it wants.

With plenty of lambdas and the ability to switch new ones on demand, we can have here what I like to call QOS in a nondestructive way. If you have some new users call up some new colors, then you map those new users to their new colors and everyone should be happy in contrast to differentiated priorities schemes, where some users get less priority than others.

A Photonic Switch

The slide below [Module Operation] shows how our wavelength selector switch works. The light comes in and strikes a diffraction grating where every color goes to a different angle. They strike micro-mirrors that reflect back on the correct angle the color that you want to have put on the output fiber. In this sense you can say I want red from the bottom fiber and blue from the top fiber to go in the middle fiber. This means then that blue from the bottom fiber cannot go in the same place because the mirror is already occupied. As a result you gain the ability to mix and match colors.

Such a photonic device costs about 10,000 dollars, obviously more if it is packaged and enabled with a controller with software. It allows flipping around some 80 different wavelengths in five or eight fibers.

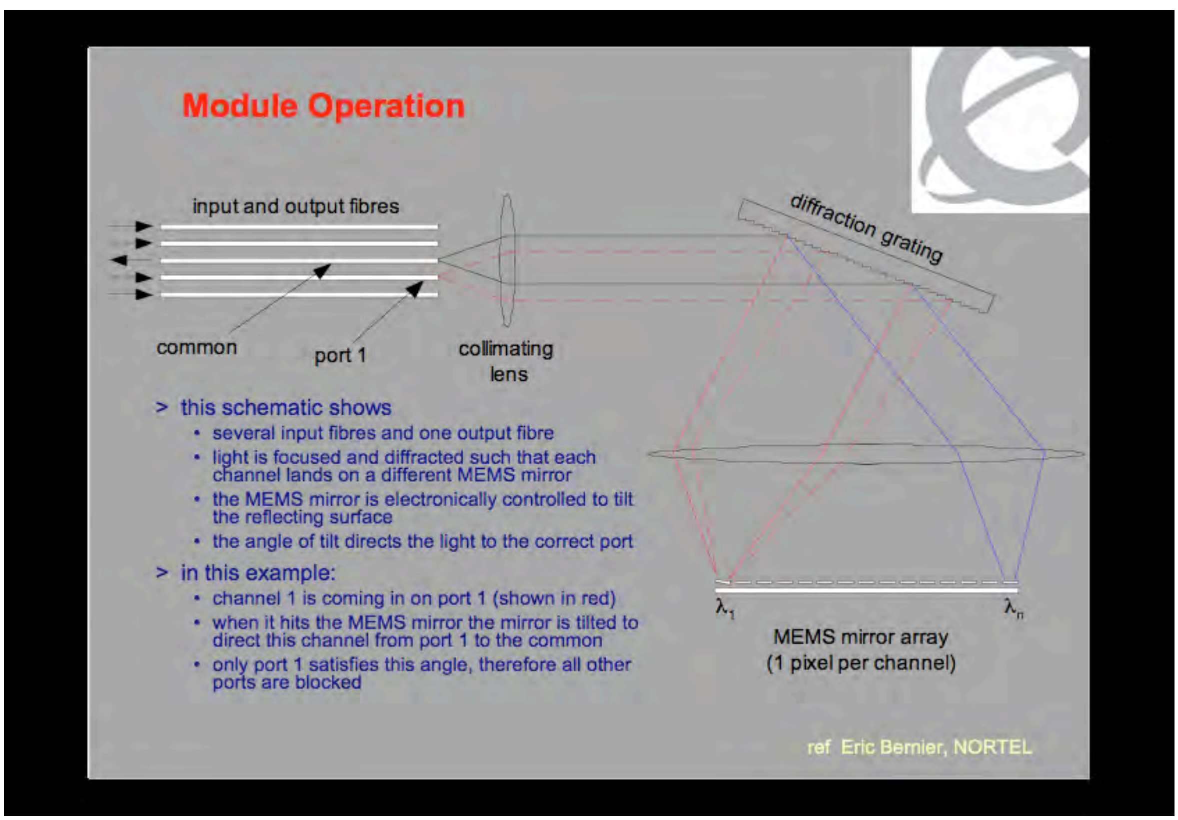

COOK Report: Is this lightpath switching?

De Laat: Yes and no. The term lightpath is used in a much broader sense. This is a very special case. This is thrue photonic switching. You are really switching colors. In other words lambdas which you mix and match. If you build an architecture composed of a number of these switches, you can send colors anywhere you want them to go without having the traffic to go to the electric domain.

Now if you want to do OEO that is optical to electrical conversion and back to optical you need colored lasers that cost you many 10ths of thousands per laser. That is very expensive. A WSS device that can do hundreds of colors photonic switching costs you also 10ths of thousands but you only need one. This is what makes photonic operation cheap.

COOK Report: Take me up a couple of levels. Here you are down at the very bottom ground floor. The purpose of doing this is that once you have done that you can mix and match colors and get your photons where you want to send them much less expensively?

De Laat: Yes, think about what we have shown in this demonstration where you play around with these topologies. You optimize the topology of the network for where you see the most load. If you see that you need a lot of capacity to a certain site you create a tree structure of lambdas to go to that site. And if in the next round you see that more next neighbor communication is needed, you make a ring topology. You can do this by flipping the mirrors on-the-fly to change the underlying topology of your photonic network. Since you do this in photonics you do not need in OEO device sitting around to make the changes.

COOK Report: But in the case of the module operation slide above, we are talking about a lambda as an optical Lightpath that can be redirected on the fly? How is this not the same as lightpath switching?

De Laat: If you call a lightpath a true lambda, then it is the same. But people also call VLANs, from ingress to egress, "lightpaths." In this sense people pollute the name "lightpath." It is not light and it's hardly a path.

Below is a slide describing how the dispersion compensating modem from Nortel operates. But this may be too technical for this discussion. Basically one figures out what the distortion of the signal will be after transport through the fiber. Then you pre-distort your signal in the opposite way such that when that pre-distorded signal gets distorted through transportation it is normal again. You

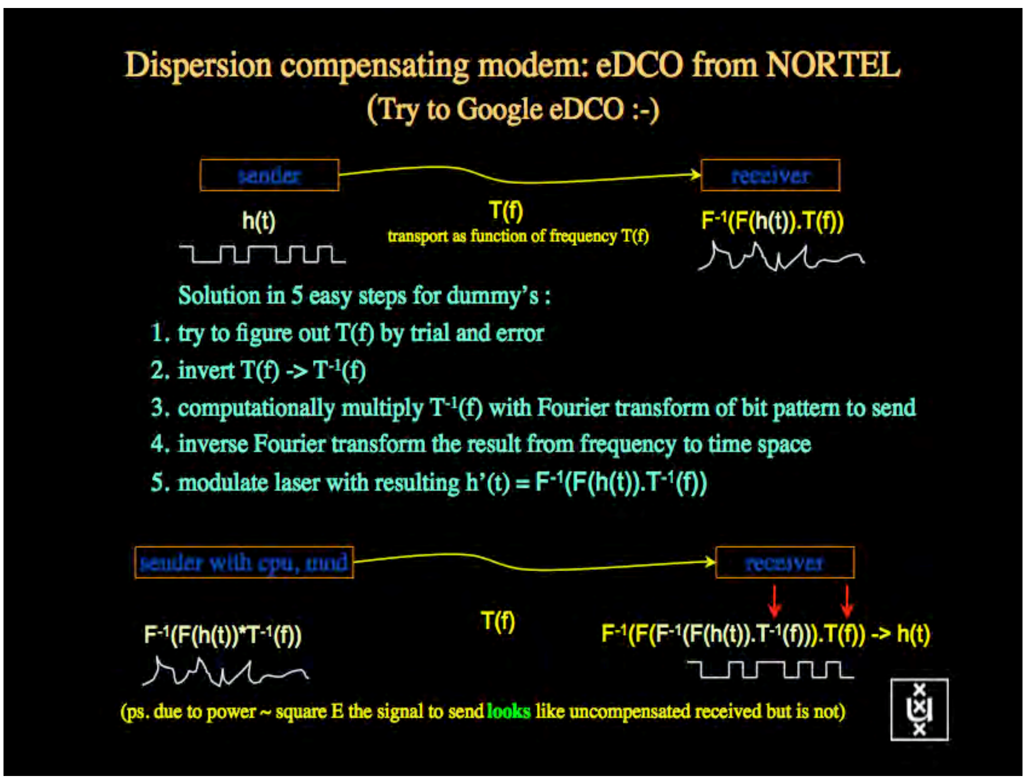

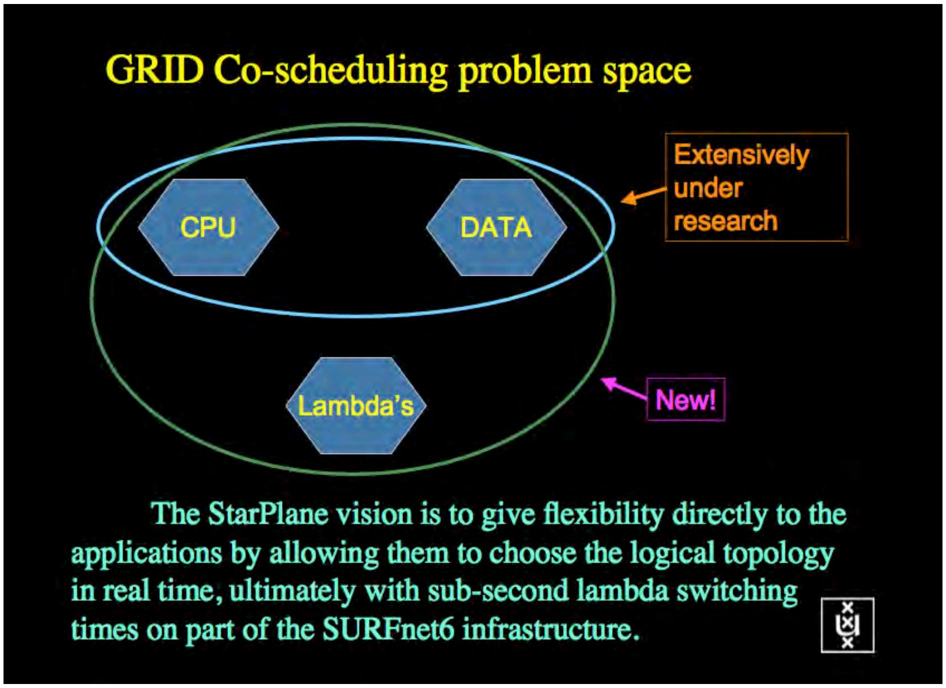

pre-distort your signal in such a way that the fiber's distortions will cancel out the pre-distortion.

Within grids you usually co-schedule compute and data resources but now you can also co-schedule your lambdas within the grid. This gives you a new playing field that is shown in the GRID Co-Scheduling Problem Space slide above. You get very constant behavior over the lamdas. That is actually trivial because there is nothing in the path which can make it non constant. Your round-trip times are always exactly the same.

Resource Description Framework Language

One of the major problems when dealing with multi layer networks spanning many domains is the description of the topologies for pathfinding. The RDF network description language is there to master and describe the kinds of infrastructure sets that we, for example, see from SciNet here at SC08 in Austin (TX). The example is from SC06.

We created an ontology to describe the components that you find in these kinds of networks.

Based on this ontology, we have schemas to describe the placement of routers, switches, interfaces, links and fibers and the framing that is used to transport data. Given that we have this ontology, we can use Web 2.0 technology to reason about it - just like friends of a friend's network where you say Johnny is a friend of Carla and Carla knows George and therefore there is a path from Johnny to George. You can take these relationships and turn them in a kind of map.

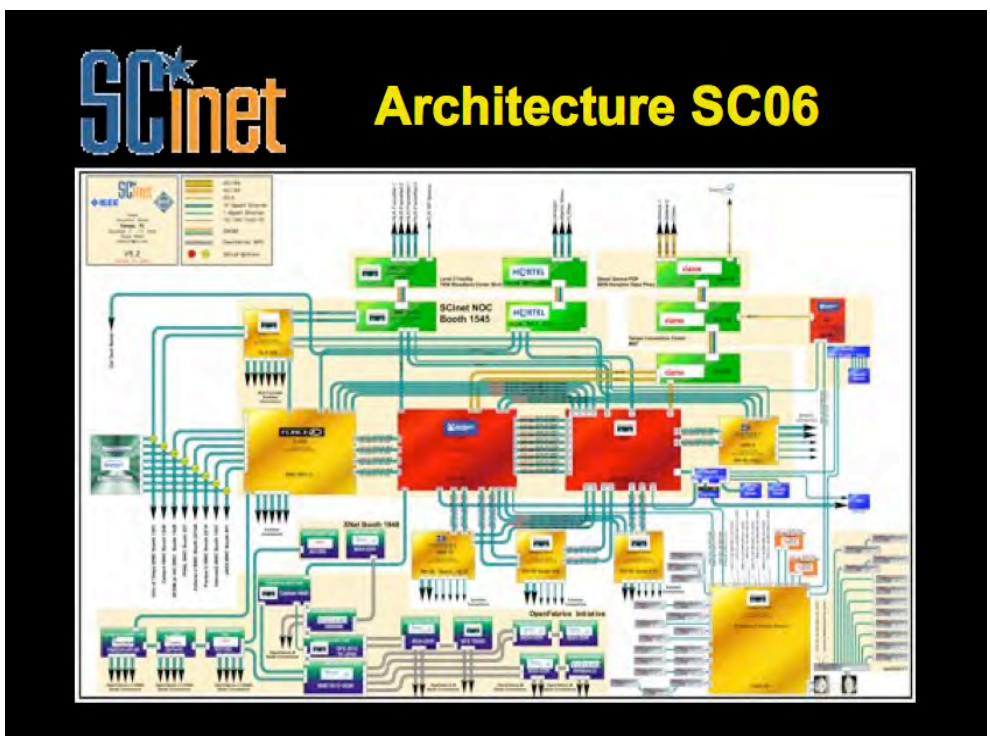

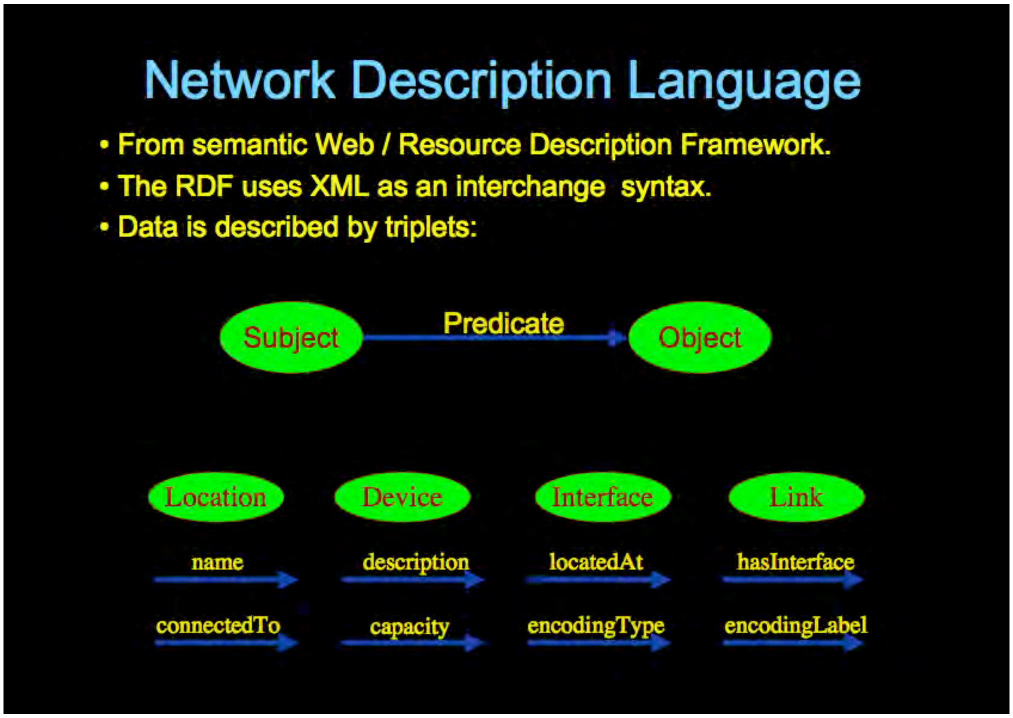

We figured out that you could do the same in networks and use Web technology to describe a friend of a friend in triplet descriptions.

The statement Carla is a friend of George is equivalent to this interface is connected to that interface by a fiber.

This slide at the bottom of this page describes the modeling process for an Ethernet and a fiber layer.

These descriptions permit us to point to other domains. Consequently you can have a description of your network and point to a description of neighbor networks, thus creating a web of descriptions. Each network maintains the description of their own network and point to the others where they connect.

We do not need one "master-of-all networks" description. You can just interlink it as a Web object. This allows you to do path finding for these descriptions.

COOK Report: And conceptually this is down in the engine room of the path finding process for which the light table on the show floor represents the graphical user interface?

De Laat: Yes. We have tools for the network description language. We showcased that by describing SURFnet in resource description framework slide at the top of the next page.

COOK Report: An NDL file is what?

De Laat: Network Description Language of which this

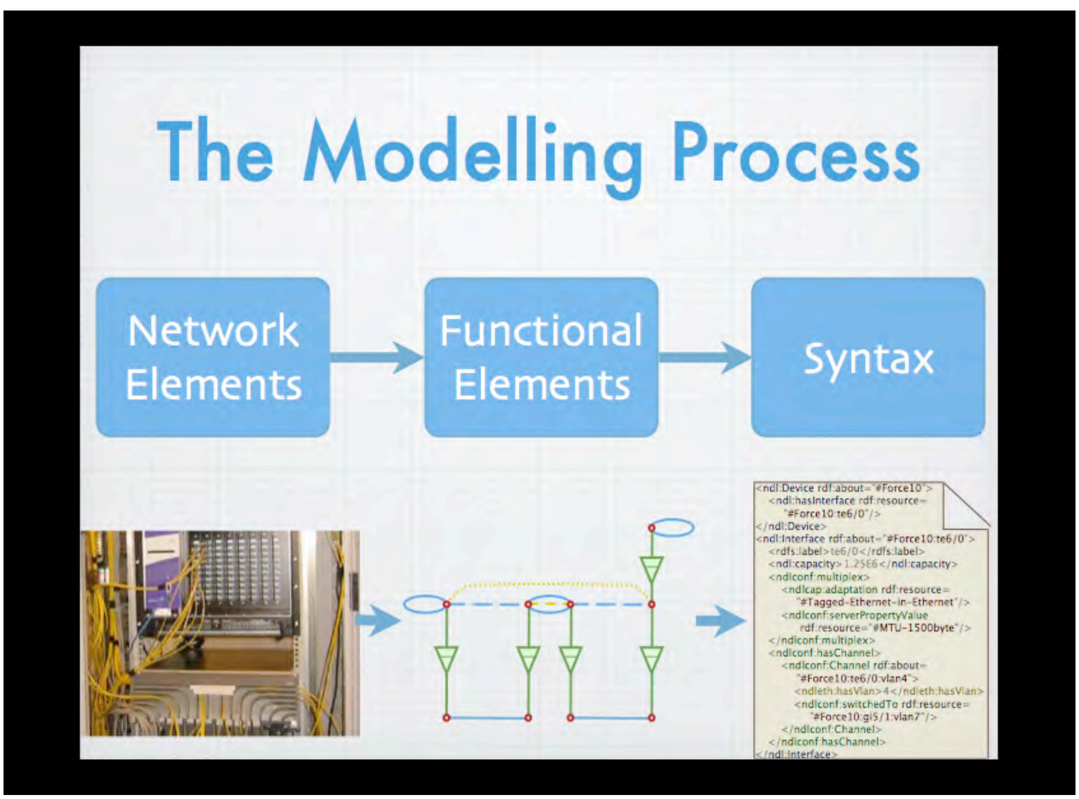

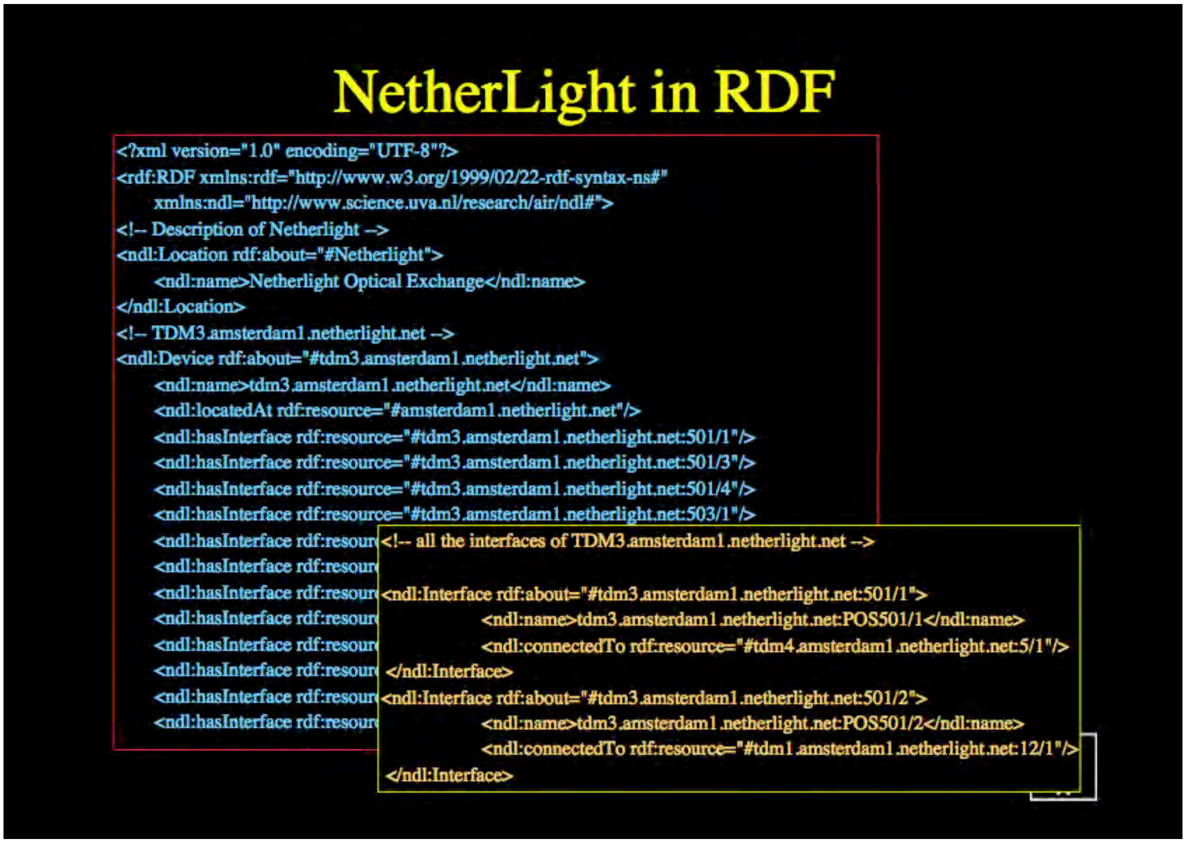

is an example.

And in the slide above presents tools to generate a NDL file for a setup and validate the syntactical correctness of an NDL file. The slide below shows an example of the multilayer features that can be described in NDL.

The slide below is a visualization of SURFnet6 based on a NDL descrip-

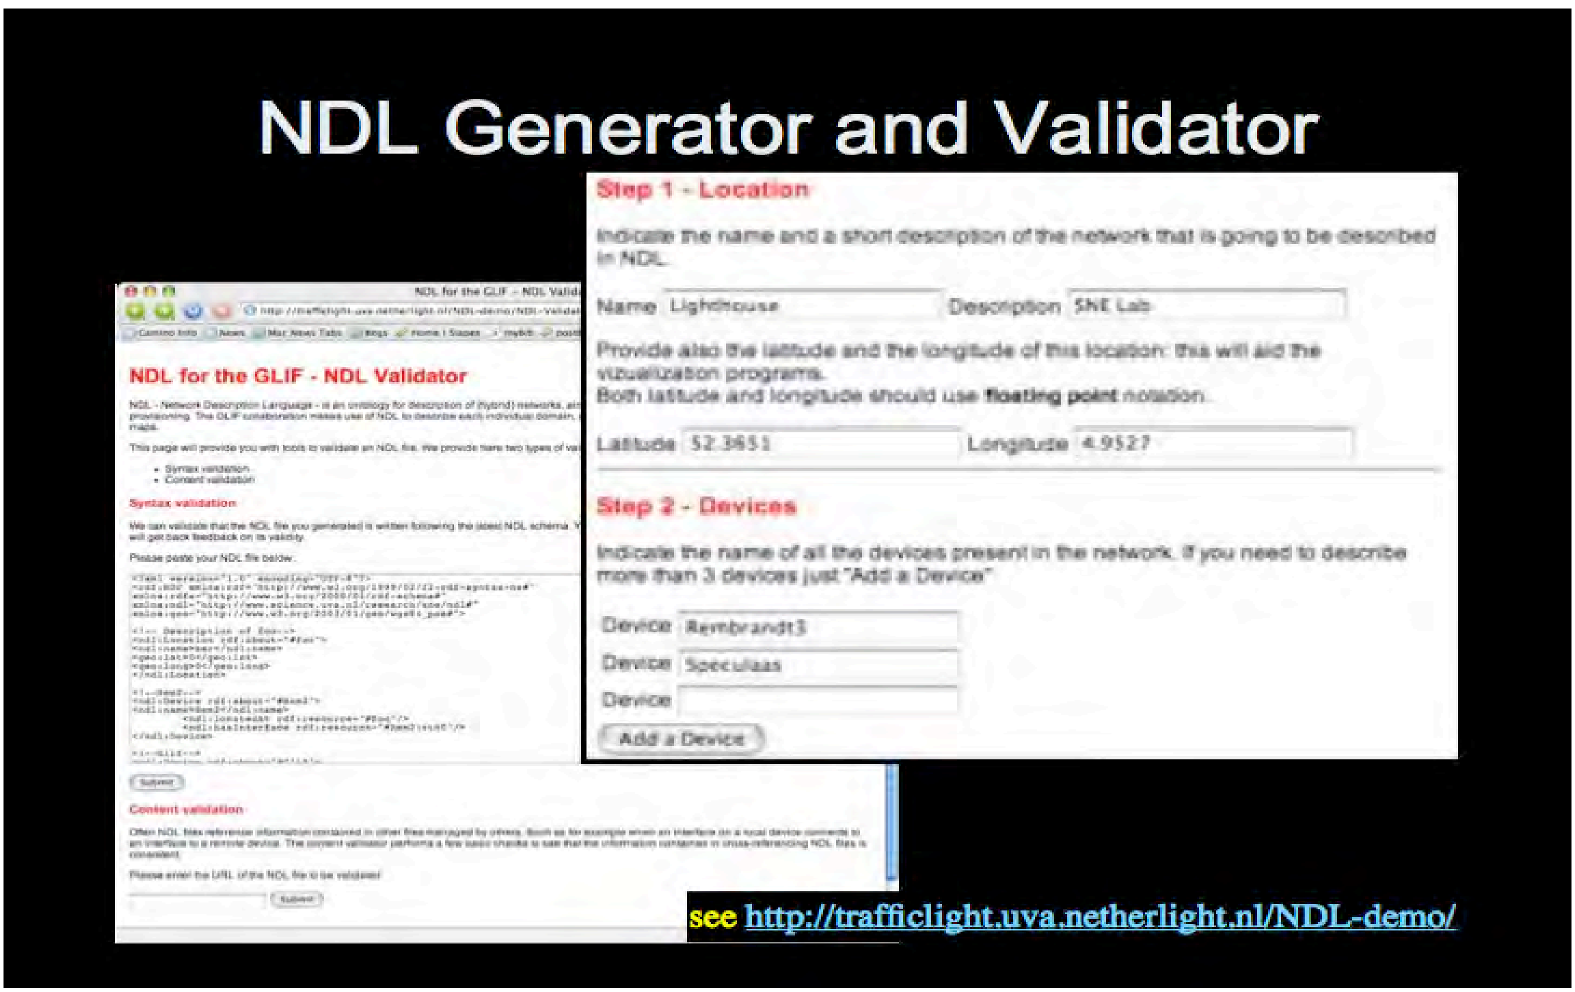

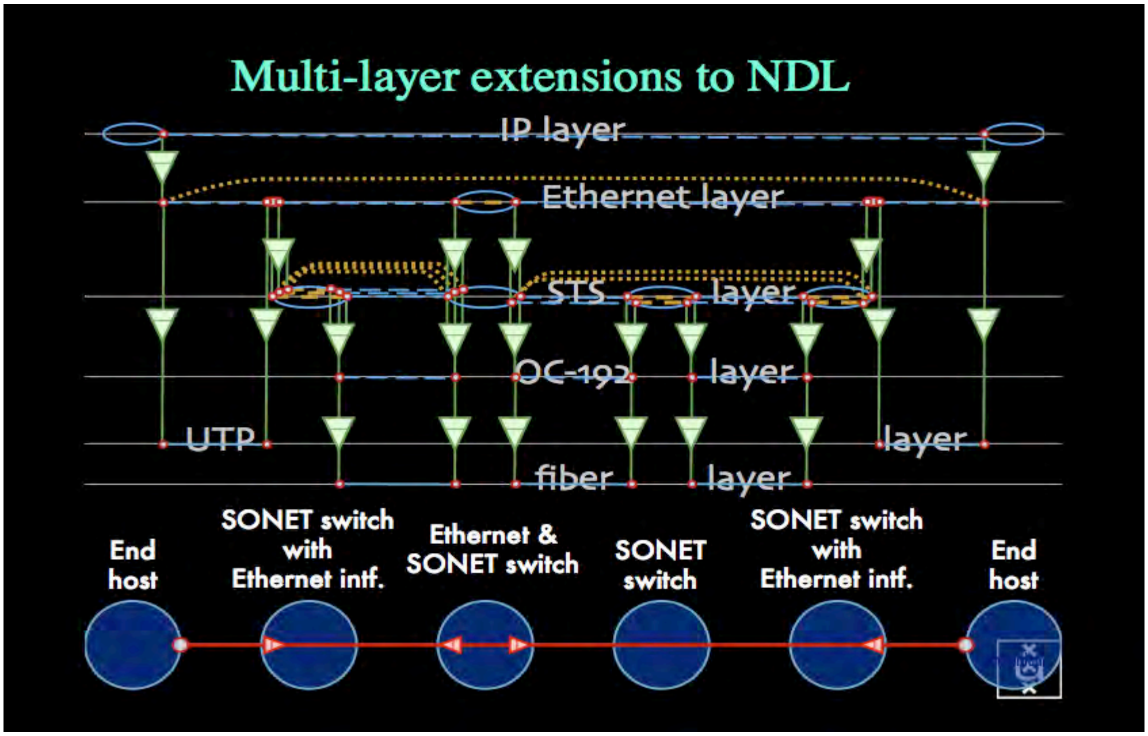

tion. Although the visualization has lots of room for improvement, it gives confidence that the model works.

Network Description Language as Routing Protocol

In the next two slides we dem-

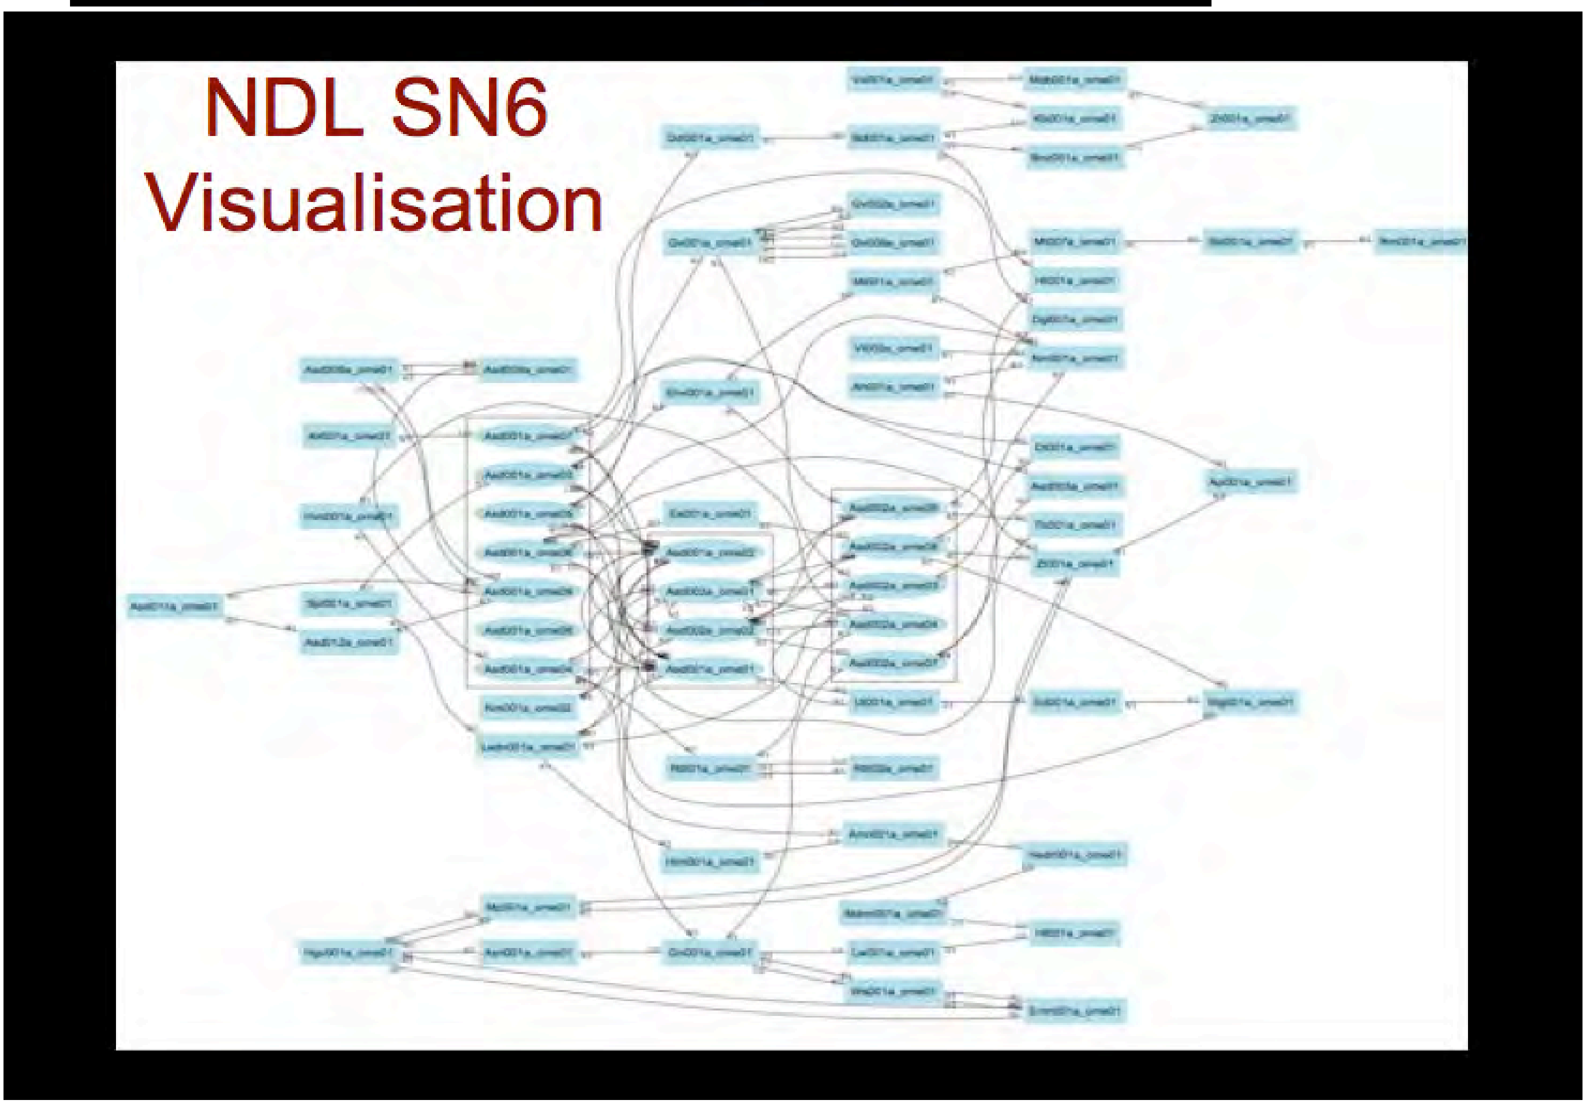

onstrate the multi domain multi layer routing power of the model in a mock-up network that requires a loop to allocate a gigabit lightpath from source to destination.

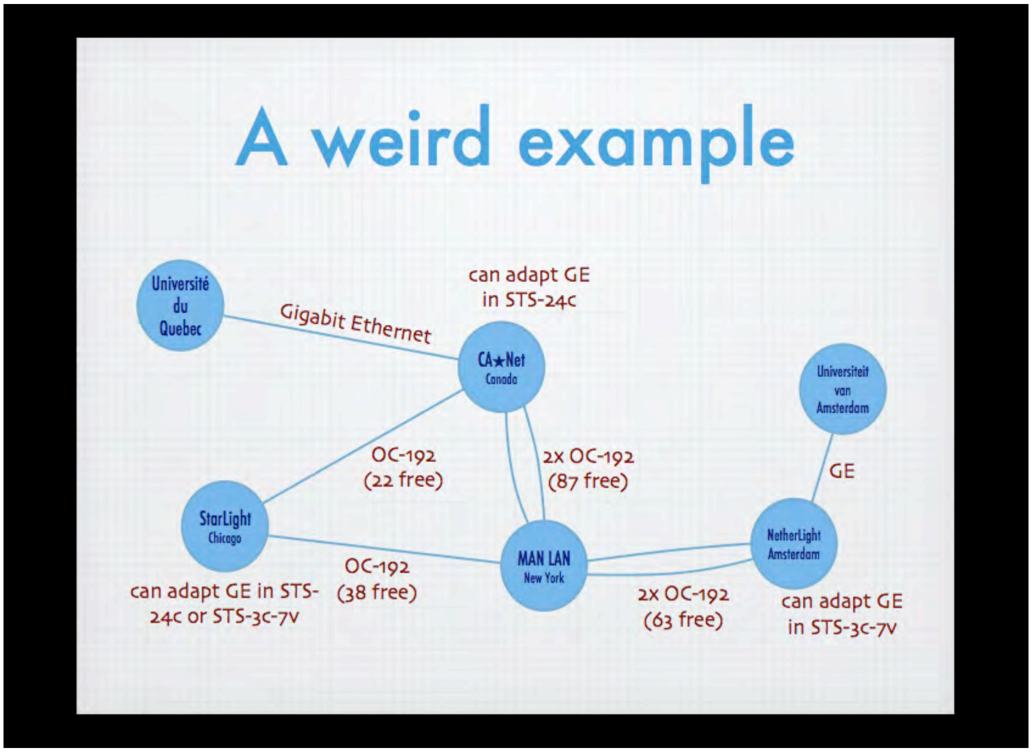

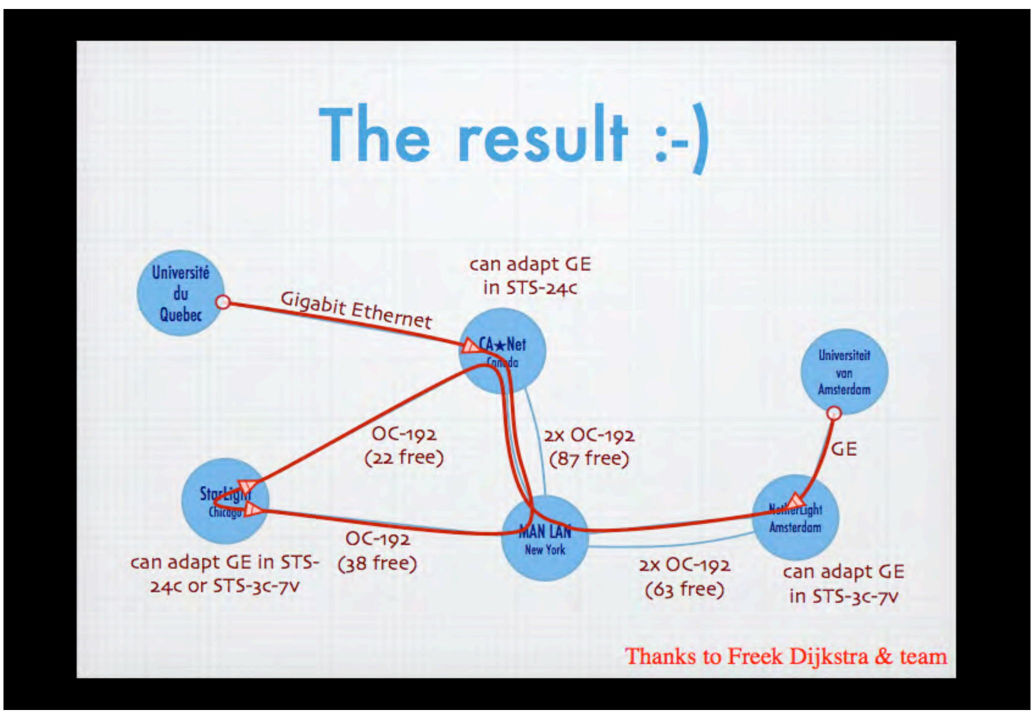

COOK Report: So in the first slide below you can not go directly from Québec to Ca-Net and Ca-Net to StarLight and then via MANLAN to Amsterdam?

De Laat: That is correct. The capacity of CA-Net to StarLight in this mock-up example is insufficient. So the traffic has to travel via MANLAN. It must go to StarLight to be translated in the correct data-framing that the destination can understand. But when it gets to StarLigth it cannot travel back to MANLAN since the remaining capacity is now also too low, but since the new framing occupies a bit less capacity the traffic can now flow to CA-Net and then to Amsterdam via MANLAN. Hence pathfinding results in the loop as shown in the slide below.

COOK Report: This is like network state information?

De Laat: Yes and it can be described in NDL that is Network Description Language. I think that there is no other operational routing protocol in the world that would permit you to be able to figure out a loop like this.. But we can do this.

And on the top of the next page is a demonstration where the protocol goes through the steps of the state machine.

By tracing the course of the Green line from the left and the right you can see how the network protocol tries all possibilities and finds its way through the maze. We are

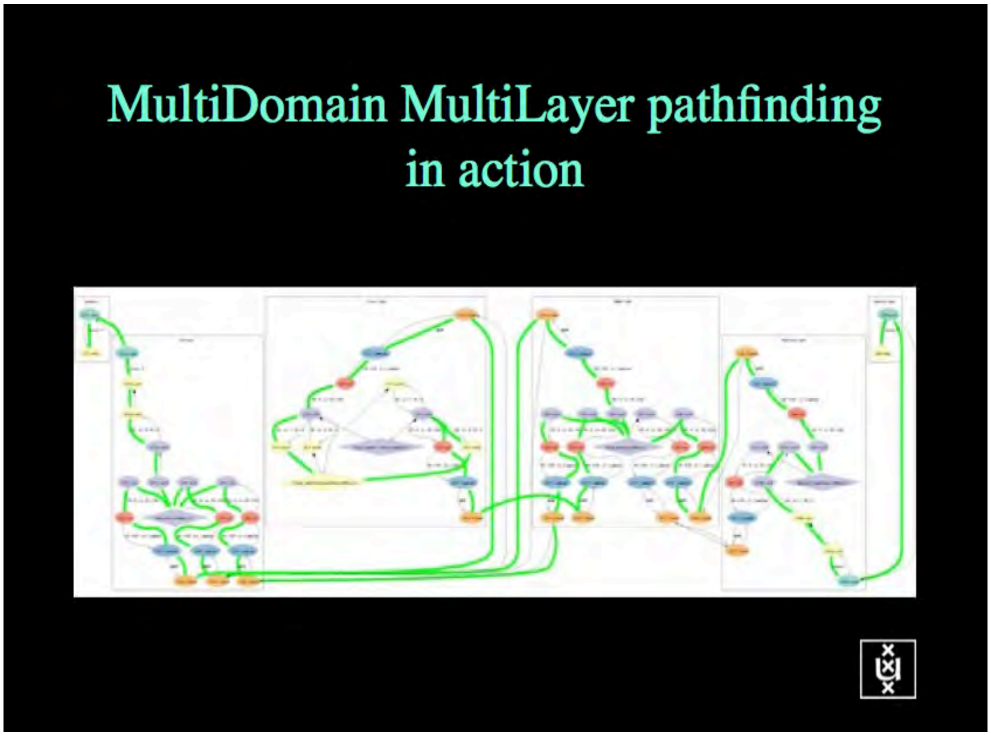

submitting NDL to standardization organizations.

We want to extend NDL to describe other types of infrastructure. If you are able to create ontologies and schemas that describe storage, content, projectors, tiled displays and transcoding services, then we are able to ask a high-level question such as: "show me video with this content and show it to me on my tiled display."

If you then find out that the video is resident on this server, that it according to its metadata needs to be transcoded to be viewable on a tiled display and what network capacities are needed, then we can request the necessary lightpaths and resources to make it happen. Given the relationships in RDF we are assured that the resources are compatible. It is just a matter of pathfinding in semantic space that translates to co-allocation of compatible resources in real space.

COOK Report: Who else in the world is doing anything like this?

De Laat: Not many groups and those that work on NML are basically working with us. The RDF wikipedia entry says "The Resource Description Framework (RDF) is a family of World Wide Web Consortium (W3C) specifications, originally designed as a metadata data model, which has come to be used as a general method of modeling information through a variety of syntax formats."

http://en.wikipedia.org/wiki/Resource_Description_Framework

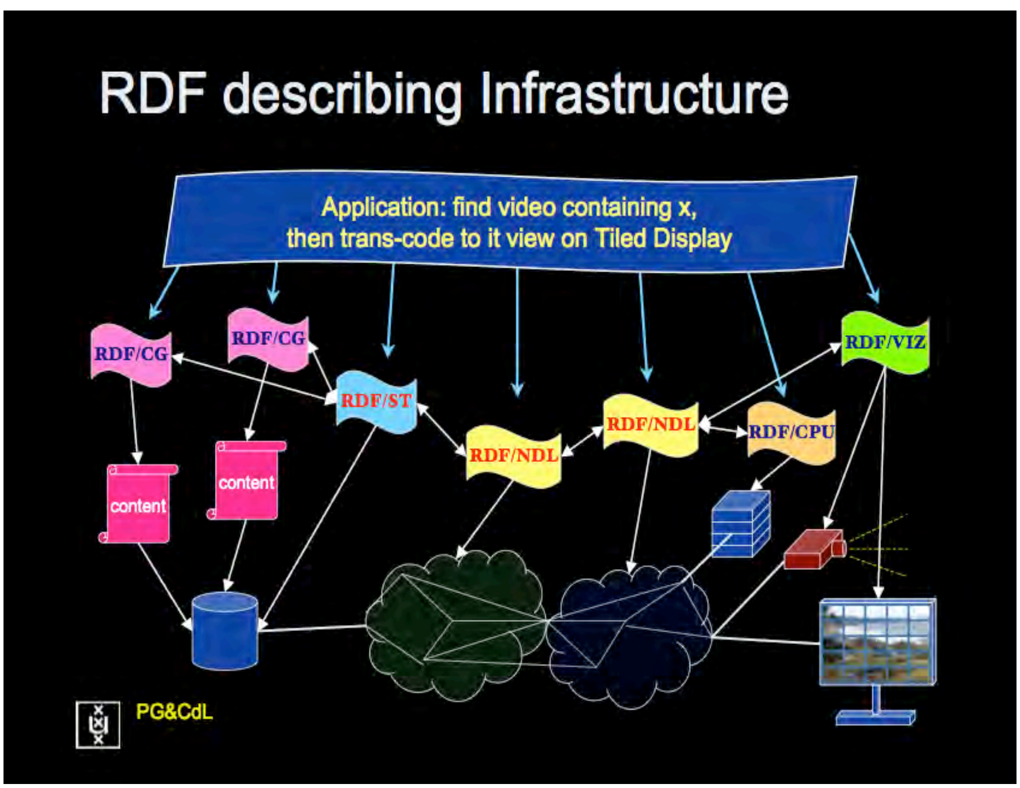

Terabit Networking

You may ask yourself what constitutes a terabit per second network? A laboratory such as CALIT2 has 8000 Gb drops. Does that constitute an 8 Tbit/s LAN? At the University of Amsterdam we have 2000 1 Gbit/s drops. Does that make a two Tbit/s lan? I don't think so because it depends on what you add up. If you look at an 64 core Intel processor and you cut it in two, the left and right half communicate at 8 Tbs per second! Of course some of the capacity on those chips is there just to let the cores talk with each other.

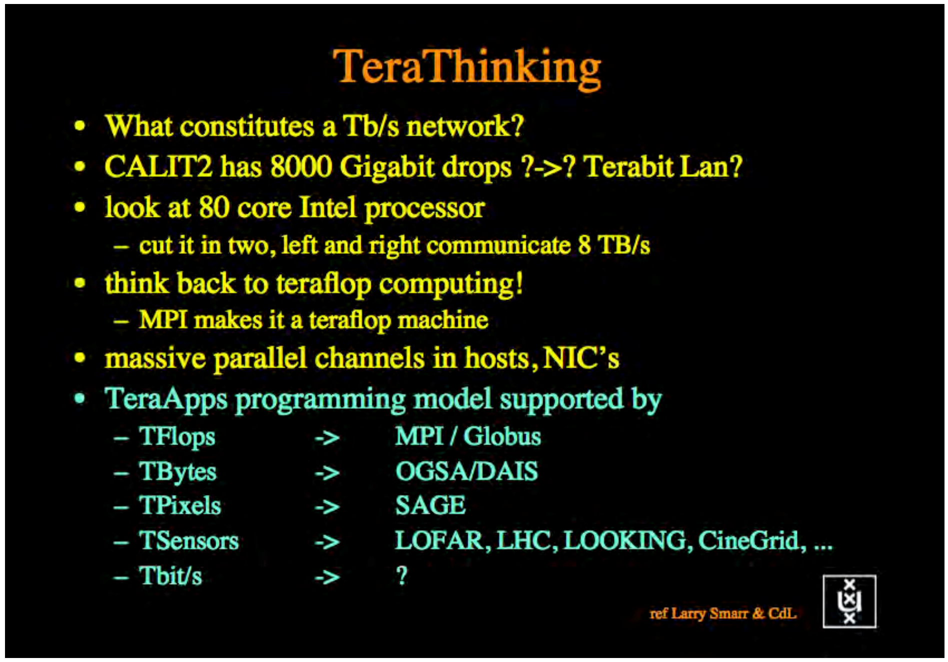

COOK Report: Your point is it's all a matter of the context in which you do your thinking?

De Laat: Correct. I consider it to be a Tbit/s network if all that capacity can be used by one application for its purposes. For that one needs programmable networks. So

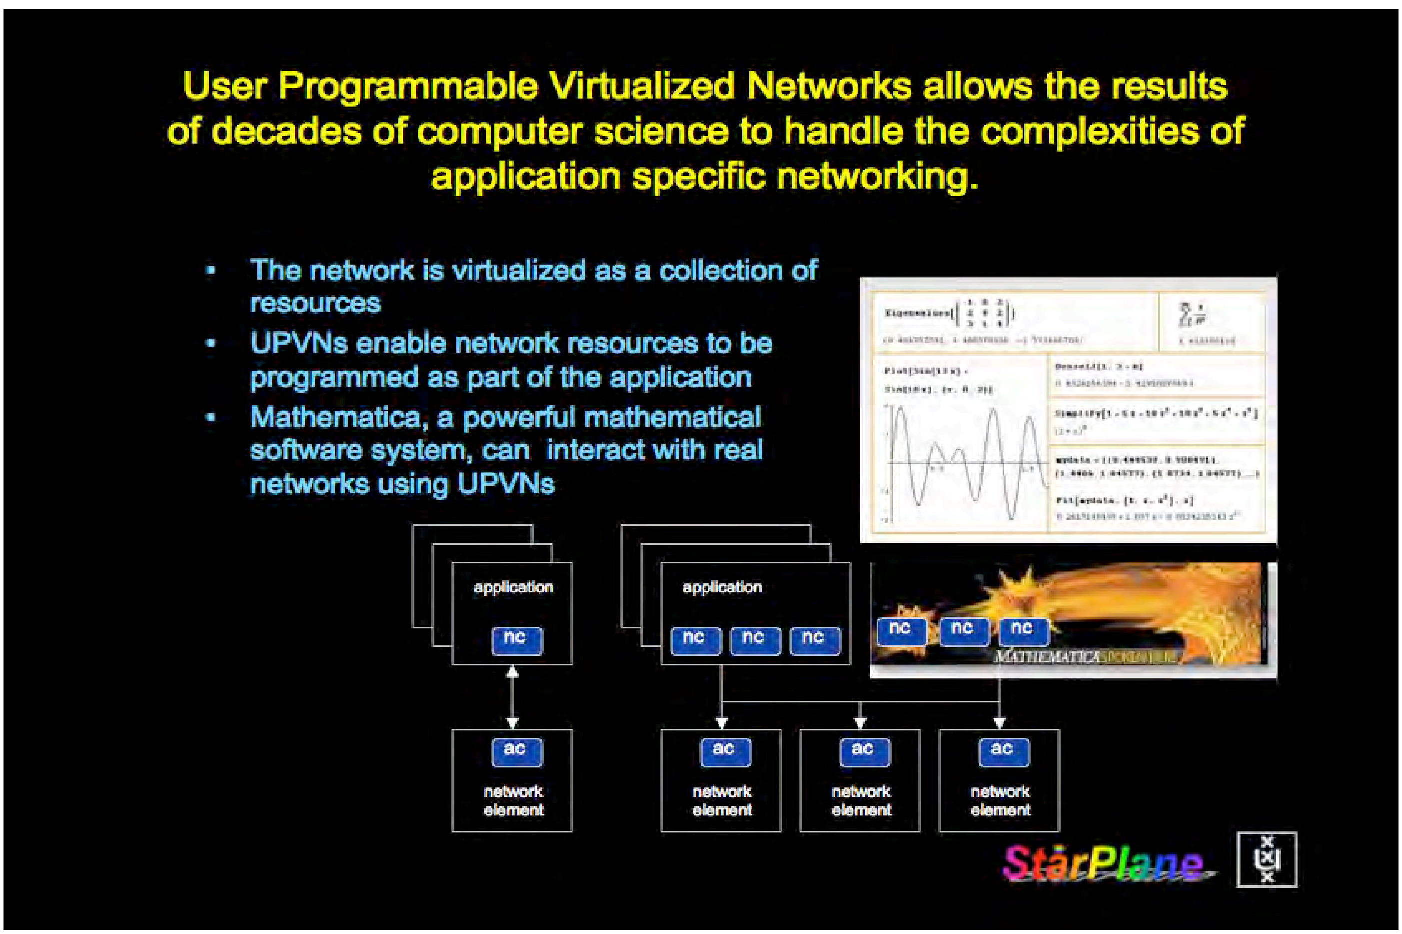

we asked the question of what constitutes a terabit network in a different way. Think back to teraflop computing. People started out here by filling up rooms full of PCs and interconnecting them. But doing this doesn't make it a teraflop computer.

What makes it teraflop computing is MPI or Globus. It is the middleware. Such middleware is necessary to ensure that an application like fluid dynamics modeling can harvest 8 teraflops worth of computing power for its own purposes. You need to have middleware to make this technology work for you at the teraflop or terabit thinking level.

To say it in a different way; the displays for an OptiPortal tiled panel are useless without the middleware that makes the tiles work as one big display.

What we are doing with our resource description framework is to drive the network to do something collective on behalf of your application. The collectivity makes it a Tb resource for the benefit of your application. But to achieve this collectivity you need programmability and middleware.

You look at it as though it were a bunch of small Lego building blocks. You say if I put this object in here I need bandwidth to talk to the core and then I must be able to write it to a disk array. You need to be able to coordinate and allocate discrete components so they work together in an organized and cooperative way.

The next slide explains the programmability of the networks by embedding the network elements in functions that can be used in Mathematica formulas.

The use of Mathematica allows to optimize the network on the fly for the problem it has to solve; i.e. getting the data at the compute elements in time for processing.

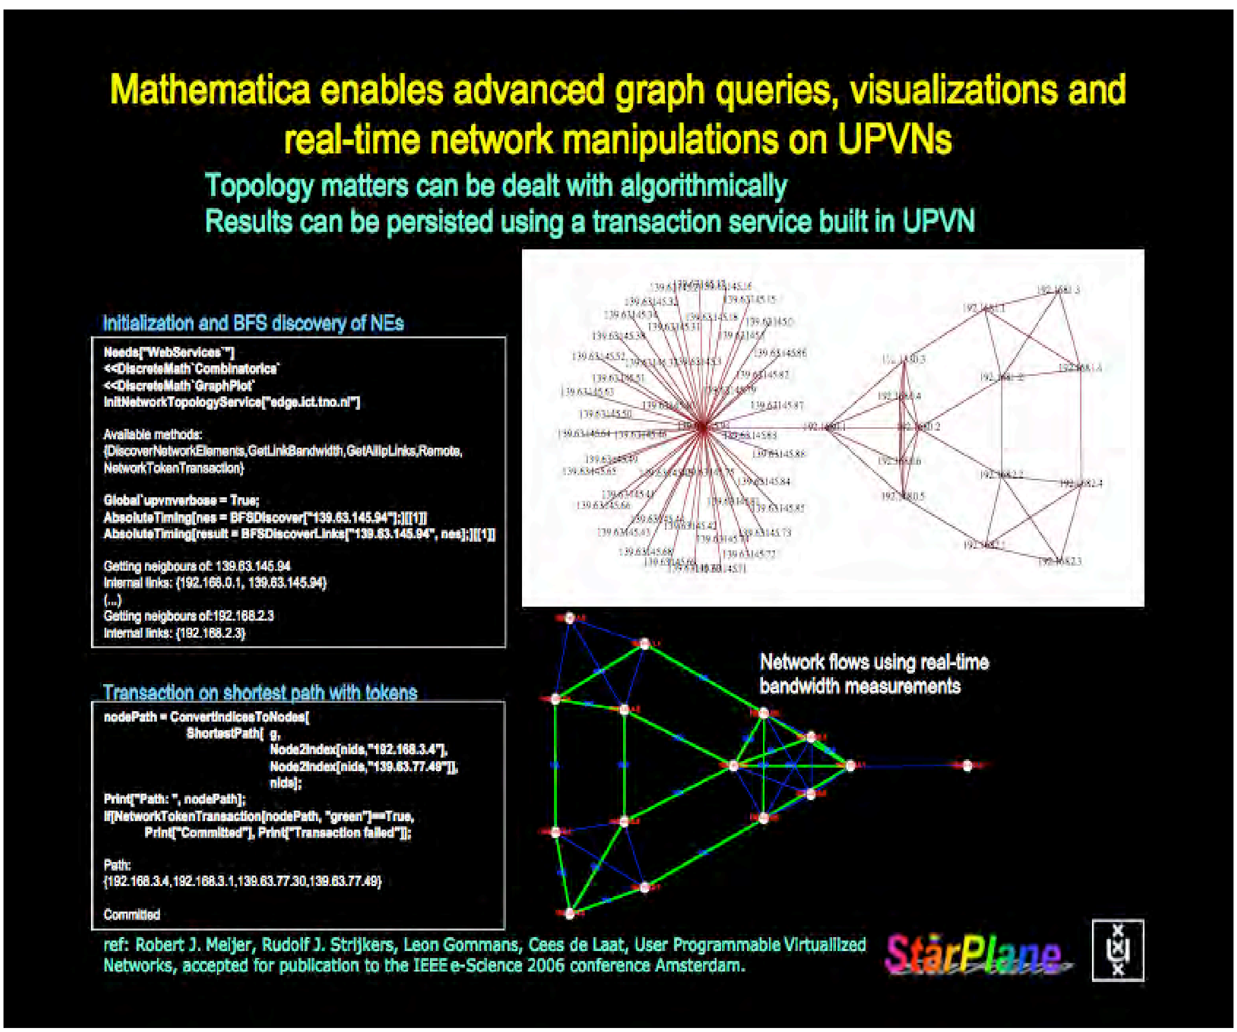

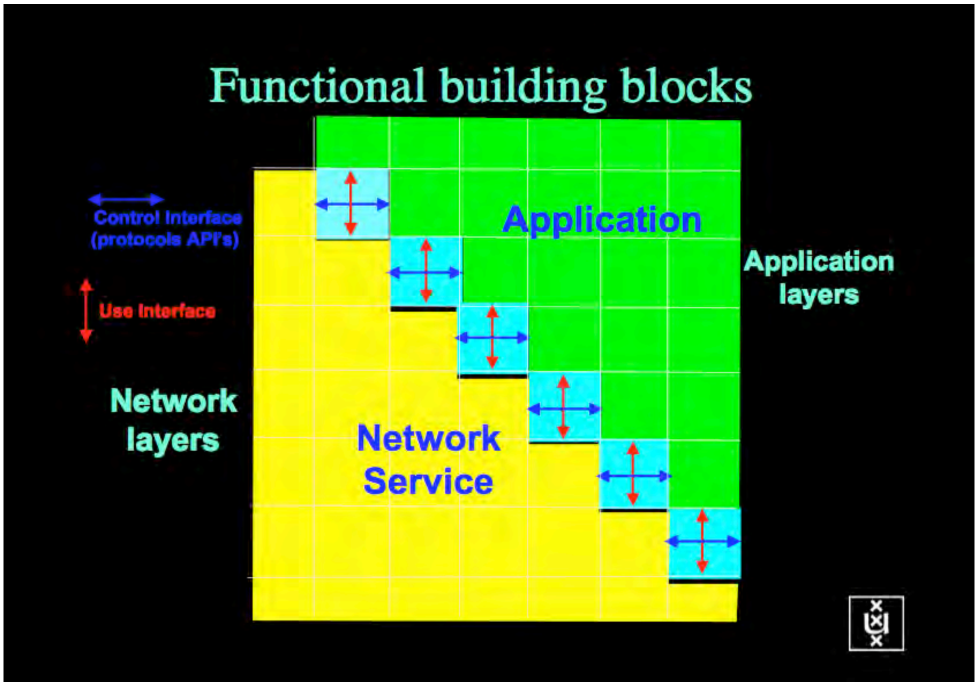

The slide above shows that applications of the future need to be aware of all the layers in the network.

Finally, in dealing with ICT infrastructure, power consumption is a big issue. The more switching you can do in the photonic domain, the more power you can save.

For every dollar of equipment you spend you can count on spending about one dollar for electricity to power the equipment during its life time. That is true for computing. Routers are so expensive that the power is relatively a minor cost factor. Still, if you put 80 colors into a router you need an expensive router interface for each color and each of those interfaces consumes 200 Watts.

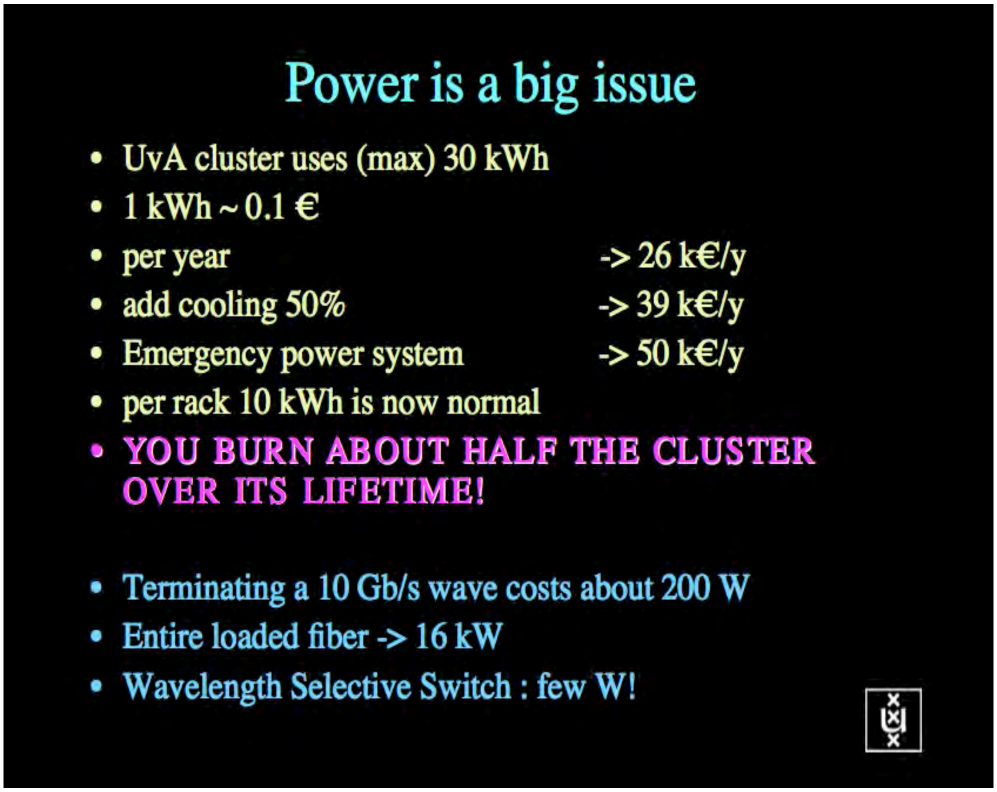

Cees De Laat at SC08 in Austin

Interactive Networks

A User Interface for Application Owned Lightpath Networks

Editor's Note: The prototype discussed below was developed by Rudolf Strijkers under supervision of Prof. Dr. Robert Meijer as part of his phd research on next-generation Internet architecture. Members of the team that developed the demo were Mihai Cristea (post-doc), Laurence Muller (scientific programmer) and Robert Belleman (head of UvA visualization lab). My Interview with Rudolf was conducted on November 19.

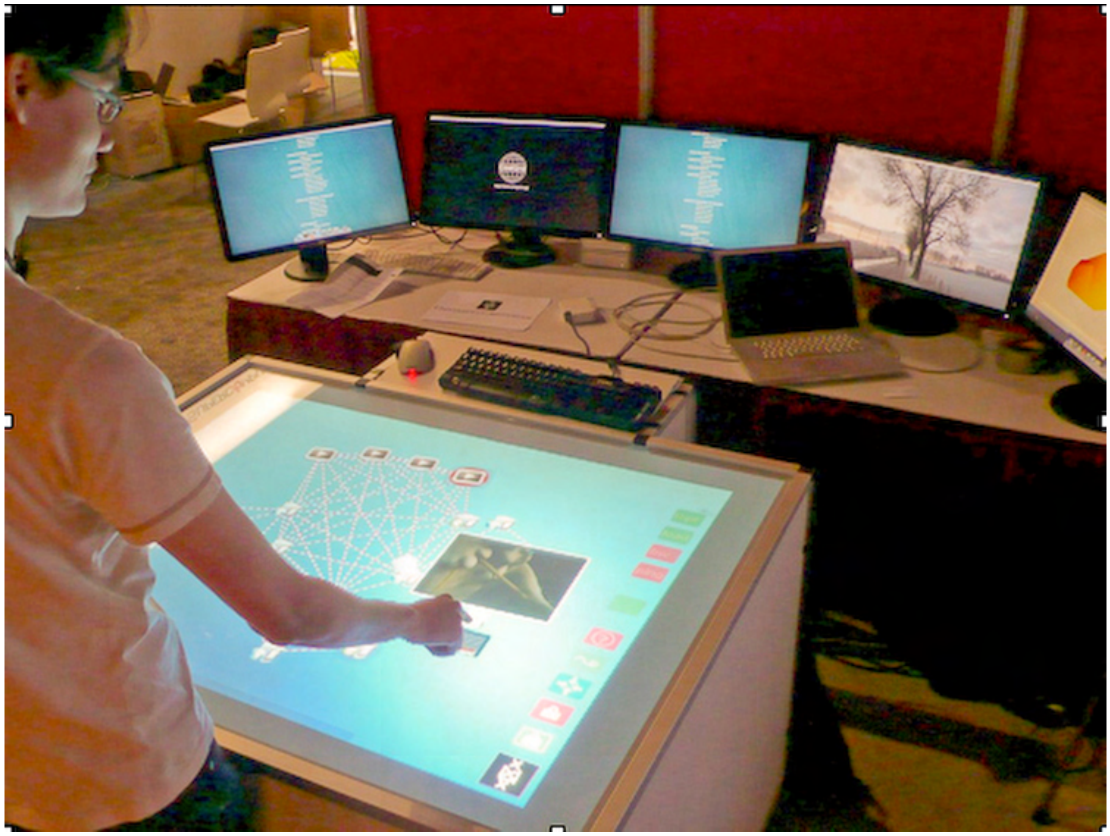

Figure 1. Interactive Networks setup at Super Computing 2008, Austin, TX.

COOK Report: Cees de Laat and I've been down in the engine room talking about what makes these user controlled and application controlled switched light path networks possible. Now we are going to look at a prototype of a user interface with a multi-touch-sensitive screen that allows a user to tap on the tools he wishes to select and with his finger to draw the paths he wishes to activate nodes on a programmable network. Multiple researchers can use the multi touch interface at the same time.

According to what you were telling me the goal of what we are talking about is to have this kind of software on an ordinary researcher's screen three or four or five years from now. At that point the user should be able to use it to control most any application in collaborative environment. Would you take me on a guided tour of what it is and what it does and how it works?

Strijkers: This is the first prototype of what we are calling "Interactive Networks". In interactive networks humans become an integral part of the control system to manage the next-generation of programmable networks and Grids. The main design principle is this; by virtualizing the configurable and programmable properties of network elements as software objects, any aspect of a network infrastructure can be manipulated from computer programs. What we show here is an implementation of an interactive control system concept for user programmable networks, which applies the architectural concepts we have developed in our research.

The network you see here is the current set up of our test bed located at the University of Amsterdam. These icons represent the network's elements and network structure visualized at IP level. I can tell you a little bit about the actual infrastructure. There are currently twenty nodes, interconnected by four subnets to create an operationally interesting topology. Three of these subnets are in separate virtual machine environments of VMware, called ESX servers, that also contain four virtual machines each. ESX is essentially a container for virtual machines. The VMware management environment enables us to create, clone, and remove virtual machines. It also enables creation and manipulation of complete virtual network environments.

The virtual subnets and machines are connected to a physical subnet, which also contains two physical nodes. Then we have the four Mac Mini's here in the booth, which are directly connected to the physical switch in Amsterdam with a gigabit connection. This way, the Mac mini's are part of the subnet.

When the network boots and the nodes come up they will connect to controller. The controller is programmed in such a way that it will send a out a discovery request when a new node connects. Each node will try to discover its neighbors using ARP to scan the whole subnet for hosts. Since it will discover everything in the subnet, we also find all the nodes in the datacenter. We currently only display the programmable nodes, but the whole discovered network can be displayed too.

At SC08 we forgot to turn off WiFi to Scinet on the Mac mini's once and discovered over 700 nodes within seconds. The neighbors you discover at Ethernet level will look like fully connected at the IP level. For example, if you interconnect the three computers with a switch, it will always look as though each computer can reach each other directly. That's why it looks like three fully connected networks here and one large fully interconnected network over there; the video screens are just one hop away.

Once we have discovered the networks you can also see the result over there in the Mathematica interface on the fifth screen. The Mathematica interface has access to the same information and supports the same network ma-

Figure 2. Network visualization after discovery. The dashed lines, when monitoring is enabled will indicate bandwidth (line width), delay (percentage of dashing) and jitter (randomize in dashing).

nipulations as the touch table. We will come back to this aspect later.

***COOK Report**:* What do these icons represent?

Strijkers: The icons represent the type of modes or functionality that each node offers. We currently have three modes. 1. A producer: this node contains streaming video content and is visualized as a green-circled arrow. It can also route traffic. 2. A consumer: such a node is connected to a streaming video client and can display the streamed content. A play button in a screen shows it. 3. A router: The sole purpose of these nodes is to route traffic and blue-circled arrows demonstrate it.

At this point we can look at what kind of videos may be streaming in the network. You can push with your finger on a producer node and a window opens on the touch table that gives you a preview of the video. By the way, the movies we currently have are: Big Buck Bunny, Elephants Dream, and two Cinegrid demo movies. The first two are actually made in the Netherlands as part of the Orange Open Movie Project and the son of Cees de Laat made one of the CineGrid movies. All the movies stream continuously stream in high definition, but to unknown destination IPs. This way a node never receives a stream, except when we run our special expressions to capture the traffic.

Underneath the node you can see a small graph, which will displays CPU load measurements. A button on the top right of the touch table will enable or disable the CPU load measurements in the network. The real-time load information for a selected node will be displayed in this graph. When the load becomes larger than 1, the nodes will light up red, alerting the operator that the node is under stress.

Now we can decide to make a path. Then we go into the path creation mode, and then you can decide to trace a path, from a producer node to a router to a screen, for example. And you see the video stream appear on the screen.

***COOK Report**:* The blue line is traced with the finger? [See Figure 4 on page 44 below.]

Strijkers: Yes. If we make a path, what will happen, we send a request to the controller asking to create the path that we just dragged with our finger. The controller will send the request to a compiler to generate the commands and forwarding expressions for provisioning the nodes. The results will be passed on to a transaction monitoring, which executes a distributed transaction to load the commands and expressions on the nodes. If loading of one of the expressions or commands fails, the whole transaction will be

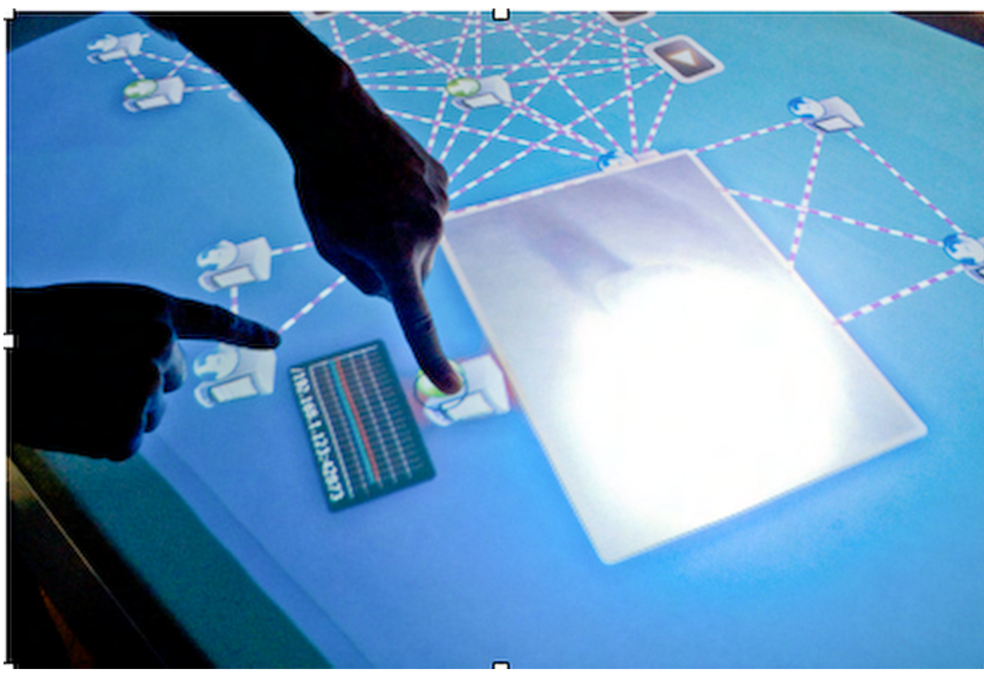

Figure 3. By touching a producer node you will see a preview of the video stream. The small graph underneath shows real-time CPU load measurements.

Figure 4. Simply drag a line with a finger to create a path.

rolled back. So, whatever happens in inserting or removing requests, the network will always be in a consistent state. When drawing a path, the touch table will also automatically select an unused color and this color will be used for lifetime of a path to identify the stream. When tracing the traffic on IP level you will actually see the color codes in the IP packets.

Once we have a path we can also select it and extend it to a multicast path. This can be done by dragging a new route starting from any node of the path chosen for the extension. It's as easy as that.

COOK Report: You are taking this content and sending it to a second screen?

Strijkers: Exactly. I only have to touch the path on a node and drag it to a screen over another route. And then you can see the movies streaming on two computers.

COOK Report: With this table, and if somebody will show me, I'm sure I will learn the basics of it pretty quick, right?

Strijkers: Yes, you will. The interface is very simple considering the capabilities. Actually, what we show here is a showcase for how we envision the management of the next-generation networks. So, the capabilities of our experimental programmable network are one step further than what you can do with modern networks now. For example, I can show you how we can draw a path with a loop. Have you ever seen a path with a loop in a network?

COOK Report: So, in other words, that path that you just set will send data back and forth two or three times?

Strijkers: Yes, the packets will ping-pong back and forth before being routed to the screen. Normally, loops in a

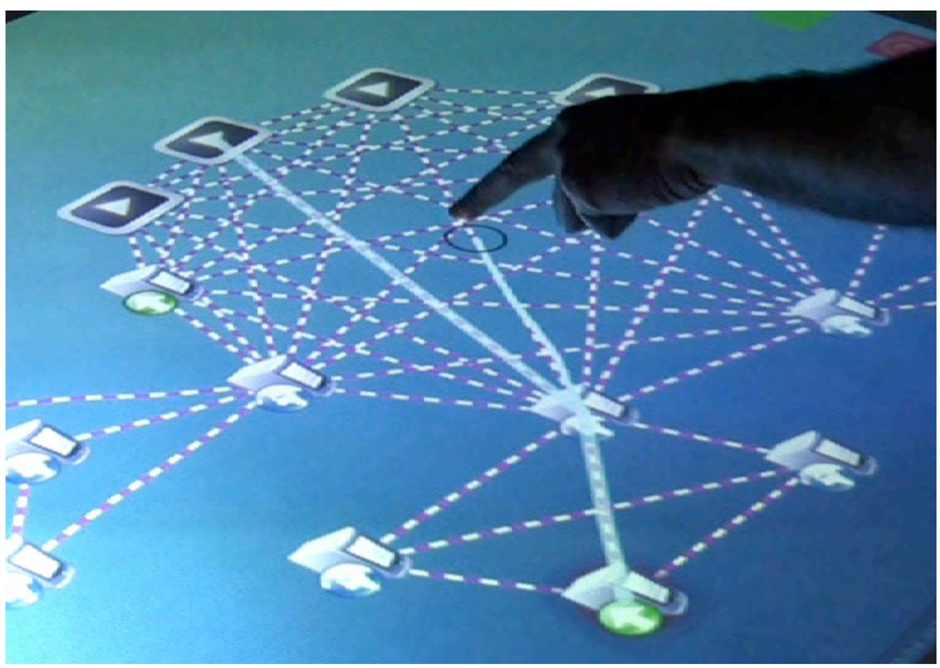

Figure 5. Selecting and extending an existing path from a node creates multicast paths.

Figure 6. Creating a loop. Unicast paths can be routed any way a user likes, whether it contains one or more loops or crossings. The compiler will detect and generate the correct expressions.

network are bad, because routers have no way of detecting if a certain packet already came by or not.

To achieve loops in a programmable network, you could use special programs with counters to detect looping packets. But, we enhanced IP a little; we put a token inside the packet, in the IP option field to be more precise. This token is not necessary, but it allows us to white list or identify packets or streams uniquely. The token enables us to bind network behaviour to traffic that is not in any sense connected to the protocols used. An example of such a binding is 'I want a good quality video connection to my TV, but only after my pizza arrives', but also to bind network behaviour of communities or distributed applications in grid networks. Our former colleague Leon Gommans did a lot of work on this subject and we have extended his ideas to programmable networks.

In the demo, the compiler detects the loop and creates expressions that change the colour of a token to a different shade. This means that at every hop the token is rewritten and will flow in another streamline graph. For example the colour still remains blue to the application, but in reality the network uses the shades to maintain state.

So, I have shown how to make a unicast path, how we made a multicast path, and how we made a path with a loop. Now I can show you a little bit what happens inside the node.

COOK Report: *OK*

Figure 7. Changing path creation mode. The button with the question mark shows the discovered topology including non-programmable nodes.

Strijkers: For this we need to switch to a different mode. This mode disables dragging of paths and allows us to interact with the nodes and edges of the network. When I double tap on a node, you will see what happened in the node when we made the loop.

COOK Report: The large black circle indicates you zoomed into the node?

Strijkers: Yes. When we make a path, like a loop for example, the request is sent to a controller. This controller runs the request through the compiler, which checks if the nodes are available, how they are connected, if the source is a video stream and if the destination is connected to a screen. The compiler will generate a flow graph for every node. This flow graph will describe how the traffic flows from the input ports to the output ports.

COOK Report: And the output port is **skb** in the red circles?

Strijkers: Yes

COOK Report: And the tbs, what does that stand for?

Strijkers: Let's start at the first filter first. Netfilter is a library in the linux kernel, which captures all the data from the networking stack at specific points. When not used, traffic would normally go through the normal networking stack of Linux. But, what we do is we capture the traffic at Netfilter input and force it to go through the flow graph. And, skb_transmit is actually the output function of the linux kernel. So, if we send a packet there, it will be routed and sent to the correct host. We have made a special modification, were we have full control over the traffic flow. This first filter is tbs. It's called the token based switch. What it does is, it looks at the packet and says, it's a blue packet, and I'm a blue graph, so I accept the packet. If I would be a red graph I wouldn't accept the blue packet. In other words, it accepts or drops packets based on their token. This allows us to create application-specific flow graphs for tokenized streams.

When the packet is accepted, it goes on to the tb and the tb filter tears off the token. Why would we tear of the token? Because if we send the message to the Mac mini, it would have no clue that we did all kind of weird stuff with the packet. It just looks like any other packet. And here is the magic; we have a filter that rewrites the IP destination of the packet to route it to go there. The library that allows us to insert/remove these filter expressions at run-time in the kernel is called Streamline. It was developed at the Vrije Universiteit of Amsterdam and our colleague Mihai Cristea and I worked closely with its developer Willem de Bruijn to make it suitable for our purpose.

Figure 8. Streamline flow graph currently loaded in a node. Note that due to the picture contrast the connecting lines are barely visible.

Figure 9. Modifying the sampling rate of a flow. The extreme left shows the manipulated stream, and the screen immediately to its right shows an unmodified stream that is also part of the same multicast tree.

COOK Report: When you get a green circle with a plus, what does that indicate?

Strijkers: The compiler automatically generates these expressions, the distributed transaction processor executes a two-phase commit on all the nodes and inputs them in Streamline. After the transaction is complete, we can zoom into the node and modify the expression that resulted from the compilation process and that is currently running in Streamline. Just by touching the plus button.

COOK Report: That gets you to a different interface, or?

Strijkers: Not exactly, it adds a filter to the run-time expression at a certain place. Keep in mind that the actual code is running in the kernel.

What actually happens is that this request goes to the specific node, it picks out the manipulated expression, plugs the sampler in and puts it back into Streamline. And you can see it, because your whole video stream goes nuts. On the left screen we see the video of a multicast branch with a sampler and on the right the unmodified stream of the other branch.

The image is distorted because we throw away some packets. Right now it throws away 50 percent of the packets. (Modifies the sampler value) So, you just saw me modifying the flow dropping rate in real-time. Now we only implemented user interface support for a sampler, but you can insert any type of filter yourself or even write your own.

This is a powerful tool to exert very fine-grained control over traffic. For example you could add or manipulate rate limiter filters, which would allow very precise traffic shaping. The operator at the touch table could manually control the traffic shaping, but it is also possible programmatically. I can show this later on.

Strijkers: And we can say, hey, we change the sampler to 80 percent. You will see that the screen gets gradually better. It's because of the encoding that it will do weird stuff. But, we can also remove it with the minus. It will go streamline again, it will remove the sampler and you will see the stream turning back to normal.

COOK Report: Impressive!

Running Mathematica

Strijkers: So, that's the touch table part. The other part is we can enable throughput measurements, so then the controller will ask to the network to continuously return throughput measurements. And you can see what happens over there.

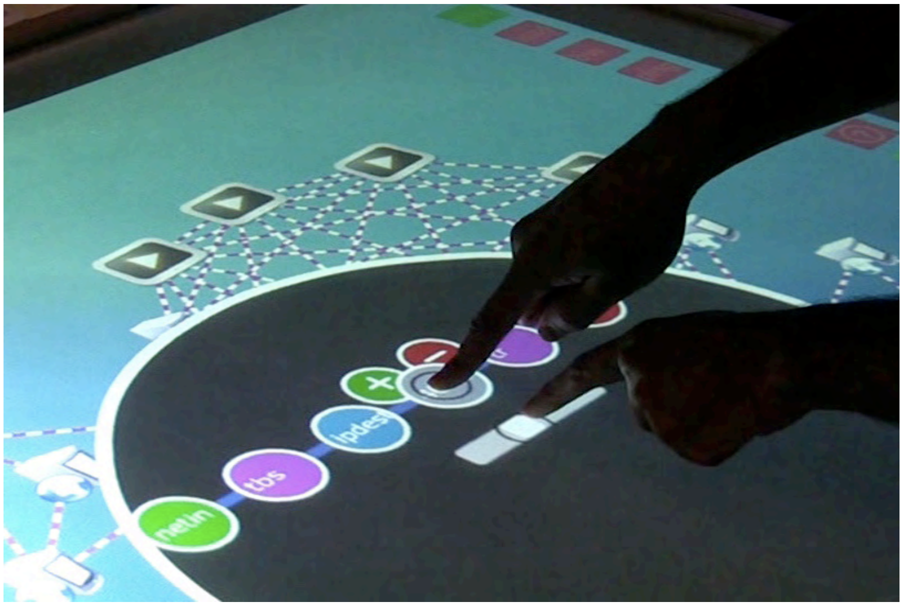

Figure 10. A closer view on the sampler modification interface. The sampling percentage is changed with the slider.

COOK Report*:* You're over at Mathematica there

Strijkers: Yeah, Mathematica is a scientific computing environment. It allows you to do interactive calculations. Mathematica contains a large library for statistical analysis, graph theory and so on. And by the way, today they released version 7, which includes out-of-the-box support for parallel computing. Mathematica also is a powerful symbolic language that allows you to write programs and dynamic visualizations. I implemented a Mathematica interface to the programmable network and now Mathematica is part of the programmable network and is able to receive measurements or manipulate paths just like the touch table.

COOK Report*:* You can see the orange yellow cone changing shape, and it looks almost like the visualisation of a pumping heart. That's my metaphor.

Strijkers: It's a 3D contour plot of the real-time throughput in the network, so the analogy with the pumping heart is quite accurate. The contour plot shows that we can now directly apply the large collection of Mathematica libraries to visualize and program networks interactively while using real-time

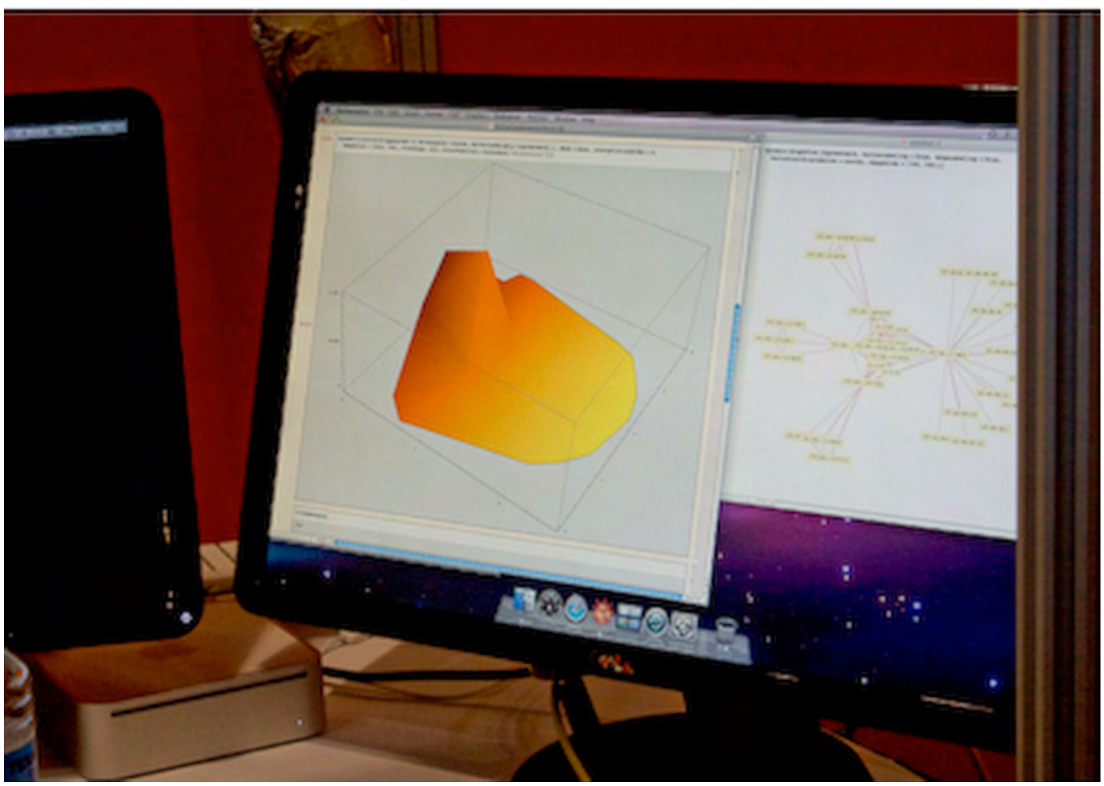

Figure 11. Interactive Networks in Mathematica. The left window shows the 3D contour plot of the real-time throughput in the network. The right window shows the current topology of our programmable test bed.

measurements. On the touch table we do graph layouts by hand for example; Mathematica can calculate the layout automatically, which you can see in the other window. To create the dynamic plot that is shown on the screen, a graph layout algorithm of Mathematica determines the form of the surface. Then the throughput measurements make up the values of the z-axis.

With the information given by the network, and currently we support continuous measurement of delay, jitter, bandwidth and throughput, one or more operators at once can write programs that automate decision-making. This opens the way for automated network adaptation in a user-friendly environment.

For example, now we can say, certain paths should avoid busy parts of the network. By using only standard functions in Mathematica, it is already possible to write a simple program that uses the real-time throughput information to continuously reroute one or more paths that avoid busy parts of the network.

COOK Report*:* In three years time, this software and capability should be on every researcher's workstation?

Strijkers: We certainly hope so. The demo we showed today illustrates a novel way to manage programmable networks. But, it is not limited to this case only. **Amongst other things, we would like to incorporate lightpath management and include other resources, such as storage or virtual machine management. We hope that on a longer term we can do trials on a larger scale, for example, by integrating our solution with Internet2, ARGIA and other research networks and Grids. Eventually, we hope to see our work applied in collaborative control room environments for monitoring and control of complex and large-scale systems.**

Conclusion

Where Is All This Headed?

The technology and projects we have just described are not destined to remain as high-end toys. With increasing improvements in hardware, software and bandwidth there is no reason that these capabilities cannot be headed for enterprise, education and consumer markets. Of course predicting any roadmap with precision is difficult. The outcome will be determined by the intelligence that underlies competing policy views.

As the interviews in this issue make clear, over the past three years, these technologies have been successfully used in international scientific experiments and collaborations. Also as Kees Neggers points out: not just for research. In the just concluded GigaPort project they have also successfully been used by ICT Departments to create optical private networks (OPNs). These OPNs link dislocations of universities and research institutes together, making centrally provided high quality e-learning facilities and collaboration tools available at all locations.

When I asked Kees what next steps he foresaw, he responded: "In the Netherlands we now have all organizations connected to SURFnet6 on SURFnet leased fiber, most of which is based on 15 year IRUs. One third of them are using lightpath services today. Since last December users can now dynamically provision and reserve lightpaths via a simple web interface. We plan to make the network even more optical and "green" and will study how we can benefit from the emerging "Next Generation Ethernet" to bring more granularity and scalability to the lightpath provisioning."

"At the application level we notice a rapid uptake by the ICT departments for their own use. Also the big science users like particle physics and astronomy, who created the first demands for lightpaths and have their own dedicated ICT know how and support, have successfully integrated lightpaths and OPNs to support the international collaborations in their fields. But outside these communities the uptake is not as good."

"We have concluded that more selective promotion and support is needed to attract users from other sectors. **ICT is too distant from the core competence of the researchers involved, and there is always a reluctance to spend money for common infrastructures with uncertain individual paybacks. Generally it is possible to get a project manager to recognize the desirability or even the necessity of the use of lightpaths but at the same time it is difficult to get it accepted on a higher level and implemented in the project execution, often because it does not fit into the original funding scheme."**

And last but not least, as proposed by ICTRegie in the report "Towards a competitive ICT infrastructure for scientific research in the Netherlands" we will have to increase the coherence between organizations, activities and the components of the ICT infrastructure.

In other words an emphasis in standardization of optical network design and project coordination among participants helps to ensure that the technology is actually used. For the researchers easily used applications tools and training and support for their use is critical. Ed Seidel as Director of the Office of Cyber Infrastructure at NSF made this point strongly in his BOF presentation in Austin saying that the High performance community in the US must focus on outreach and application tool development.

On the research end Cees de Laat pointed out that major goals include:

1. Development of common interfaces for application software to interact with the new networks so that distributed applications can use the special features of the networks. Standardization efforts are needed here and are underway in a number of standards organizations.

2. Development of a topology distribution infrastructure for pathfinding through these networks. Here also efforts are underway.

3. Greening the network by keeping traffic as long as possible in the photonic domain. This includes light passing through other domains. Technically this should be possible; culturally it is a challenge.

4. Study using simulations of the scaling properties of such networks with respect to aggregation strategies, and control and management protocols.

5. Virtualization using WEB 2.0 technology to enable the operation of the programmable network paradigm across domains.

6. Policy and authorization frameworks.

7. Interlinking different NREN hybrid networks such as Internet2, GN2/3, Phosphorus test bed, G-Lambda and others for dynamic service setup.

There is no longer anything that is hindering static service setup. So many pieces of what we developed in the past seven years can now be implemented by everyone today.

Cees also put it very well when he wrote: making this broadly available will require a shift at the provider's business models where they open up as much as possible the infrastructure and offer it to the customers. It seems that the current model is to cut up service in small slices and sell that for the highest possible price. The problem is that the infrastructure needed for that cutting up is expensive and becomes another item for which customers must pay. Some-one mentioned once to me that 90% of the bill of telephony was for the metering instrumentation and processing of that bill and 10% went in the costs of the actual call. So what can we learn from that? :-)

For Further Investigation - the OptiPuter

Two issues of *Future Generation Computer Systems* (FGCS) are worthwhile reading. The first is Volume 25 issue 2 Feb 2009 entirely devoted to the OptiPuter Global Collaboratory and *FGCS,* Volume 22, Issue Number 8, October 2006 on "iGrid 2005: The Global Lambda Integrated facility."

See also information about Optiputer found at http://www.optiputer.net/:

"The OptIPuter, so named for its use of Optical networking, Internet Protocol, computer storage, processing and visualization technologies, is an envisioned infrastructure that will tightly couple computational resources over parallel optical networks using the IP communication mechanism.

The OptIPuter exploits a new world in which the central architectural element is optical networking, not computers - creating "supernetworks". This paradigm shift requires large-scale applications-driven, system experiments and a broad multidisciplinary team to un-

derstand and develop innovative solutions for a "LambdaGrid" world. The goal of this new architecture is to enable scientists who are generating terabytes and petabytes of data to interactively visualize, analyze, and correlate their data from multiple storage sites connected to optical networks.

And then look at Optiplanet: http://www.evl.uic.edu/cavern/optiplanet/ "From 2002-2008, a global team of researchers and networking engineers has been building the OptIPuter, a National Science Foundation-funded initiative to dynamically configure a distributed computational facility, where the optical network becomes the "backplane" connecting high-end computing, storage and visualization resources." The OptiPuter has its own wiki. http://wiki.optiputer.net/optiportal/index.php/Main_Page

This is fascinating since among other things it teaches everyone how to build their own OptiPortals. Everything is open source and not stratospherically expensive.

There are approximately 40 OptiPortals installed around the world. The illustration above shows 12 of the early ones.

It's the Process and the Planning

In the process of putting together this issue on the development of lightpaths networks in the Netherlands and the creation of a virtual laboratory for e-science to further technology transfer between Dutch universities and industry, one of the things that has struck me most strongly is the very clearly defined approach taken by the Ministry of Economic Affairs, ministry of education, Science and Culture and the Dutch equivalent of the national science foundation to establish and organization known as ICTRegie.

Reading the five year plan of the Dutch ICT Research and Innovation Authority shows the development of a very sharply focused way of approaching the expenditure of government funding for research and innovation on behalf of the interests of society at-large. The plan is a 41 page pdf downloadable at http://www.ictregie.nl/index.php?pageId=6&pubId=15 (Go to this page and click on the link "downloads.")

Although I quoted this in the introduction on page 2 above, it bears repeating. The decision to form ICTRegie was based on the understanding that "**within the complex field of ICT, with its many research disciplines on the supply side, countless application domains on the demand side and a large number of organizations throughout, a purely top-down style of direction is unlikely to be effective**.

Accordingly, ICTRegie **mobilizes** stakeholders on both the supply and demand side, encouraging them to join each other in thinking about opportunities for innovation. It **challenges** them to inspire each other to arrive at promising innovations using ICT, **encourages** them to seize the opportunities for innovation, and brings those opportunities together within a national vision."

"Recognizing that real break-throughs come form cross disciplinary efforts, ICTRegie adopts a roadmap for dealing with four areas of strength in each of three disciplines. It then sets forth procedures to achieve a balanced portfolio of investments in ICT research and innovation by balanced cross disciplinary work in four market sectors, four social domains and four application areas within ICT.

Economic Domains -

Here the first domain is:

Hi-tech systems and materials - the field of high-tech systems and materials is confined to the high-tech manufacturing industry, i.e. machinery and the material-reliant bio, pharmaceutical and chemical sectors. Modern technical equipment (such as production machinery, cars, baggage handling systems and high-speed screening systems) is controlled by embedded software. . . . .

Creative Industry - In this sector, the Authority wishes to concentrate on the field of ICT & the New Media, which includes the developers of innovative content and the companies/organizations which provide innovative ICT solutions to disseminate this content as television programmes, websites, games and various multimedia applications.

Food and Flowers - The 'Food & flowers' sector involves activities such as product enhancement, horticulture, the food and confectionery industry and the raw materials supply industry. The sector is incontrovertibly a strength of the Netherlands, largely due to a finely balanced interplay between technical control of production processes, transport and storage, as well as carefully coordinated cooperation between all links in the chain from producer to consumer

Water - the greater computational capacity now available enables mathematical models with a growing degree of integration to be produced. Combined with data-gathering and interpretation techniques, this opens the way to the development of new control instruments, perhaps linking meteorology with water management, or in the form of innovative inspection methods for waterways and water barriers.

Social Domains -

Heathcare - (in three parts)

- The healthcare client: the Smart Health Surroundings theme is concerned with the design of an intelligent healthcare environment which promotes the autonomy and self-determination of the patient.
- The healthcare professional: the themes Knowledge Management and Imaging diagnostics and image-directed treatment are intended to provide the healthcare professional with the level of knowledge and information required to offer high-quality care to patients.
- The healthcare manager: the theme 'Tailor-made Care' helps the managers of healthcare institutes to cope with the complexity involved in providing customized care without incurring additional costs.

Safety and security - Safety and security involve applications in the fields of defence, peacekeeping, anti-terrorism operations, crime reduction and crisis management, as well as ensuring the continuity of vital amenities such as energy, telecommunications, public administration and logistics.

Education - All schools now have computers and broadband Internet connections. . . . In the next phase, state-of-the-art ICT applications are to be integrated with developments in educational reform. This will enable new styles of teaching and learning to be undertaken, being not only more efficient, but also more effective with a greater degree of individualization, opening up the vista of true 'lifelong learning'.

Mobility - Mobility is concerned with both passenger and freight transport using all available modalities (car, bus, train, boat, aircraft) as well as the relevant infrastructure. The accessibility of major economic centres such as cities, industrial parks, harbours, and airports remains a significant societal problem. Predictability and the opportunity to choose between the various modalities based on accurate information are important aspects. . . .

The ICT Sector -

Ambient Intelligence - The emergence of ubiquitous computing, natural man-machine interfaces, wireless networks and ever more intelligent ('smart') systems allows ICT to be applied in creating an intelligent human environment. Indeed, ICT is itself ubiquitous and is accepted as a 'given'. . . .. . ICT has become a matter of course, a part of everyday life. This is a typically European concept in which the focus is on people (from computer-centred computing to human-centred computing) and has formed the main theme of European ICT research for several years. It is also in keeping with the typical European competences.

Product Software - or industry specific software - The Authority intends to establish how general knowledge, theories and (business) strategies from the world of software engineering can be rendered more widely applicable by means of product software development. In doing so, it is important to investigate which business strategies have proven successful in the internationalization of product software, and what product and service strategy has proven effective in spreading the risks.

Enterprise Information Systems- (EIS) support the various business processes within an organization. Examples include Enterprise Resource Planning (ERP) systems, data warehouses and workflow management systems. Companies in the financial services sector, public administration and production companies alike are now having to integrate their primary processes with supporting ICT to an ever greater extent.

Services science - The first is that the nature of ICT service provision is changing. It is shifting from an isolated activity geared towards cost reduction and increased productivity to one which seeks to provide added value to clients. Integrated, distributed, service- oriented architectures, ... The second reason is technological in nature. The combination of huge computational capacity with virtually unlimited networking and storage capacity has enabled today's ICT to provide new forms of service rapidly and accessibly.

COOK Report: All of the preceding are quotes from pages 18 -- 20 of the plan. The common thread through out is that the efforts are applied to problems not in isolation, but rather from seeing the problems as complex interactive systems. The thread points out that knowledge silos cannot thrive as silos any longer - that they are inevitably linked by the growth of powerful computers ubiquitous and unlimited storage and fast, capacious networks. ICT is the common glue that binds it all together and that sustainable economic growth will be only found through application of this glue.

In adopting a systemic way of looking at innovation, ICTRegie does something that, as far as I know, has never before successfully attempted. On page 21 it states

"Innovation is a creative and dynamic process. Its essence is the communication between researchers, application developers and users which will lead to new ideas. Creativity is required to generate those new ideas, but also to arrive at new combinations of existing ideas, the 'cross-pollination' between different ways of thinking, 'outside the box' thinking and trans-sectoral innovation. In many cases, this involves a multidisciplinary process. Success relies on leadership and vision." . . . .

The aspects of innovation may be represented as follows.

- **Business opportunity**: creativity in the formulation of requirements and the concomitant 'value proposition'.
- **Business model**: creativity in the organization of the value chain in order to achieve the proposed added value for the domain in question.
- **Technology**: creativity in the manner in which the technology is applied in fulfilling the requirements of the client/user, with due regard for functionality, ergonomics, quality and ease of implementation.

"Mobilizing, inspiring and encouraging

ICTRegie has opted for a style of direction which facilitates rather than dictates. The Authority wishes to mobilize the parties on both the supply and demand sides to take part in the Communities of Interest and to consider opportunities for innovation together. It challenges them to inspire each other towards promising ICT innovations, and it encourages them to seize and elaborate the opportunities thus identified.

In addition, the ICTRegie Advisory Council will produce an annual strategic report with its recommendations to promote focus. The report will be public and will serve as the basis for debate. The report, together with the recommendations of the Communities of Interest will form the input for the National ICT innovation vision and ICTRegie's annual working plan." p. 24

COOK Report: What the ICTRegie Plan seems to be attempting to institutionalize is the cooperative, collaborative nature of innovation on the internet. Can we create a framework for problem solving where our best minds will solve problems in a way that benefits Dutch society?

The American Approach Needs To Move in the Dutch Direction

I can't help but notice and be somewhat discouraged by the contrasting American approach – an approach that has worked well when we had a capital base large enough to support it, but one that in the context of economic meltdown is likely no longer sustainable. Here are for the last few decades we have tended to make large appropriations to agencies like the National Science Foundation and to measure progress by the percentage with which we can increase those appropriations every year. The process is one of throwing money at problems.

President-elect Obama has pledged to create a national CTO. It seems to me that ICTRegie should serve as a role model for the definition and function for this new and very important position.

Now since Kees Neggers pointed out to me that Ed Seidel is both a man who 'gets it' and, since assuming the directorship of the Office of Cyber Infrastructure at NSF on September 1 2008 is in a position to do something about our problems, I have spent some time within the Office of Cyber Infrastructure portion of the NSF website looking at proposals, programs, and papers. The 2003 Atkins report that established the office is quite broad and sweeping. Exactly how we are to get from here to there is not nearly as well defined as it is in the ICTRegie five-year plan.

The 2007 OCI vision document – the most recent that I could find – is detailed from a discipline point of view. The process is more general and it seems to me that the process at this point is more important. Publications on the OCI website seemed to reflect a sizable cross-section of University research interests without being tied down in any way that reflects with any precision how they will contribute to economic growth and development in the United States.

While this is in part explainable by the mission of the National Science Foundation to support basic research, one must wonder whether a mission in basic research isolated from an understanding of its impact on national economic productivity in our country will be sustainable in the context of our recent losses of several trillions of dollars as a result of the unregulated Wall Street binge?

Nevertheless based on a brief meeting with Ed at SC08 and on discussions with others I conclude that he understands the issues discussed here should be given enthusiastic support as he undertakes this important new mission.

Some Thoughts on What is Needed

In my opinion SURFnet is the leading optical network in the world and the SURFnet builders have found that while they have created many useful tools, researchers and enterprise scientists in Holland do not use these tools in their everyday work and, at a minimum, need training in understanding how the tools can be a productive use of their time.

But these tools involve far more than just Lightpaths, they involve grid software, Web 2.0 services and highly trained technical people who are only now beginning to develop collaborative research programs in interdisciplinary fields like those the Dutch have pioneered at their GigaPort over the last five years.

What I saw at the Supercomputer meeting in Austin holds huge promise and is very exciting. The booths were filled with what would look like magic, to the non-technical person. But those who create the magic realize that they must begin to develop a program to transfer what they do to the commercial and educational world.

The Dutch have been working on a small scale with Phillips, IBM and Unilever for several years. They understand what has to be done. Managing these tools to facilitate technology transfer is a complex task that requires deliberate and careful planning.

But my experience also has shown that these tools are real, extremely powerful and very much deserving of careful dissemination. After all, the interviews herein show how SURFnet has, at least since iGrid2002, been trying to very carefully nurture the process that has just come to completion with the end of the Gigaport project today. (December 31, 2008.)

In short the ***COOK Report*** concludes that the technology works and that the most significant variable will be the speed of dissemination. The speed and resulting benefits to national economies will be determined by the ability of governments to invest in national fiber networks as infrastructure that extends the benefits of ICT primarily to the society at large rather than to carrier shareholders.

If I were President Obama what then would I do?

1. Figure out how to acquire for cash or for other "suasion" a federal IRU on fiber pairs on Level 3's, ATT and Verizon's national footprints.
2. Give federal dispensation for the RONs (Regional Optical Networks) to join their fiber into those pairs forming a National Innovation Network.
3. Offer encouragement for enterprises to participate - long term - in e-science on this innovation network with an understanding that nothing miraculous will happen over night.
4. Elevate the role the NSF Office of Cyber infrastructure to a much more public and permanent focus as the focal point for enterprise outreach education coordination.

5. The FCC, the new National CTO, and the directors of Obama's infrastructure implementation team should let all the states know that fiber interconnection with their nearest RON is strongly encouraged.
6. There should be some significant education efforts undertaken by the administration's media people - via and similar internet tools YouTube - to explain to Americans why this is worth while. This "stuff" is intensely Visual and if you have had no immersion in it, the impact is easy to dismiss.
7. The cost of doing this is not HUGE, the payback immense. after the crash of the ponzi economy it would be nice to have some vision to inspire students that there was something more meaningful than Wall Street in their future. (See Paul Krugman, The Madoff Economy" http://www.nytimes.com/2008/12/19/opinion/19krugman.html?scp=1&sq=krugman%20ponzi&st=cse)

These are highways of light that will subsume all forms of communication education, commerce and knowledge transmission. The sooner they become ubiquitous, the more broadly their educational and collaborative efforts in all forms of science - especially those requiring modeling on a large scale can be invested into the building of our human capital.

A knowledgeable friend complained that requiring any national fiber network to unbundle one or more fiber pairs was tantamount to forcing it to commit suicide and to causing the implosion of the telecommunications industry, putting additional huge numbers of out of work. Faced with this concern, I must ask whether we subsidized the horse and buggy whip industry at the expense of cars because we were afraid of unemployment?

I submit that this opinion is exaggerated and assert besides that **we cannot afford not to wisely invest what economic resources we have left.**

Be it a new ARPA or an American ICTRegie our new administration MUST grasp the present opportunity to sweep out the wreckage of cronyism and speculative capital. In the national and public interest like Roosevelt did in the 1930s, it must build out a "rural electrification system" - this time an interstate highway system of optical fiber.

I end with a reminder of Harvey Newman's quote of a month ago.

"The focus on video as the motivation for true broadband [must be] temporary."

"Network applications involving access to, and sharing of large volumes of binary data as the basis of information, and ultimately as a basis of knowledge, are highly developed, but are not so visible in the world of entertainment and social networking, as they are in the realm of research.

But soon corporations will learn to follow in the footsteps of the research community to handle and benefit from the knowledge implicit in such datasets, whether for healthcare or for other business processes, or for new forms of education, that complement web-page and video (more traditional) 'content'."

"Even in the days when walls of your home are live displays (the walls themselves, as extensions of current OLED developments, not just screens), it will be the knowledge behind the images, and the ways they are used to inform and educate, as well as entertain, that will matter most."

And just before Christmas Harvey sent me a private note which he has given me permission to publish. I was pointing him to JP Rangaswamis' blog saying

This man ***gets*** it!! He is pushing BT toward becoming a services based cloud computing platform.

Harvey Newman: I agree with this. And it makes me think - Often, to make one'sn in-cumbent, one is constrained to say "What is the business model of the Internet ?"

In this case what they are asking is: How do companies ("how do I") make a profit out of this ?

So I say -

Only the highway-building model ("building the nation from the bottom up") will get us out of the past.

And if we do not do this - there is strategic as well as economic risk. Since any nation can do this."

I will be further exploring the implications of this via an interview with Harvey in the March issue.

Appendix: E-Science Recommendations

ICTRegie

Netherlands ICT Research and Innovation Authority

6.8. Users of the ICT Infrastructure

SARA and NCF have recently published 42 interviews with professors at Dutch universities in a large variety of scientific fields using the Huygens supercomputer. Each of them points out the importance of high performance computing for his or her field, the societal relevance and impact and their actual Petascale challenge.[25] More than 100 research groups rely on the use of the national supercomputer Huygens and the national compute cluster LISA, and publish high quality papers in high-ranked scientific journals. Every Dutch university and research center has an excellent connectivity through SURFnet6. Many important scientific disciplines like high energy physics, astronomy and life sciences (NBIC) have early adapted grid technologies and are now using the BiG Grid infrastructure. Other disciplines like cognitive sciences, medical sciences and media are growing. The use of the ICT infrastructure in the social sciences and humanities is emerging. The Virtual Knowledge Studio KNAW aims to support researchers from these disciplines in the Netherlands in the creation of new scholarly practices, and to stimulate reflection on e-research. A core feature of the VKS is the integration of analysis and design in a close cooperation between social scientists, humanities researchers, information technology experts and information scientists. This integrated approach should provide insights in the e-research that can contribute to new research questions and methods in the humanities and social sciences. The VKS collaborates with the Erasmus University Rotterdam in the Erasmus Virtual Knowledge Studio KNAW (in short: Erasmus Studio) based in Rotterdam and with the Maastricht University in the Maastricht Virtual Knowledge Studio KNAW based in Maastricht.[26]

6.9. Relevance for industry

Apart from the direct interest of a few high-end and large industrial research entities (like Philips and DSM), industry is most interested in the expertise and knowledge generated by the availability of an advanced ICT infrastructure. It allows them to apply and use this expertise in their own research labs and develop their own infrastructures (e.g. Shell). The ICT industry (equipment, software and services) has a similar interest in the growing base of expertise, but also in the opportunity to contribute to innovative developments and the participation in pilots and test beds.

Considering this industry position, it seems that there are no realistic opportunities for Public Private Partnerships in operating the proposed advanced infrastructure. For the development activities, the proven model of public tenders, as profitably and successfully used by GigaPort/SURFnet should be continued, as well as the industry cooperation in research consortia like in VL-e.

7. Recommendations

As motivated in the earlier sections, achieving the desired top position as a knowledge economy asks for competing on scientific and economic markets. Maintaining the position of the Netherlands in the top of science requires permanent investments and developments in computers, databases and tools of the ICT infrastructure.

Today, the Netherlands has a very good starting point: a world leading position in the field of communication networks, a competitive position in supercomputing (Huygens) and leading edge knowledge development in de field of grid technologies and e-science. The Netherlands can use this position to further strengthen it in order to attract talented scientists pushing the limits even further and to attract international companies for locating their research and development in the Netherlands.

To make this possible, it is required to increase the coherence between organizations, activities and the components of the ICT infrastructure. The ICT infrastructure must be used by a wide range of scientific fields, including alpha and gamma sciences. The activities of the VL-e project need to be continued and extended. The funding mechanisms must be changed, and become independent of existing mechanisms for investment and research projects. The ambition is to bring all components of the full ICT infrastructure at the level of the current world leading position for communication services. To achieve this goal, ICTRegie has formulated the following recommendations.

7.1. Governance of the ICT infrastructure

Position all ICT infrastructure development and operations under the SURF umbrella.

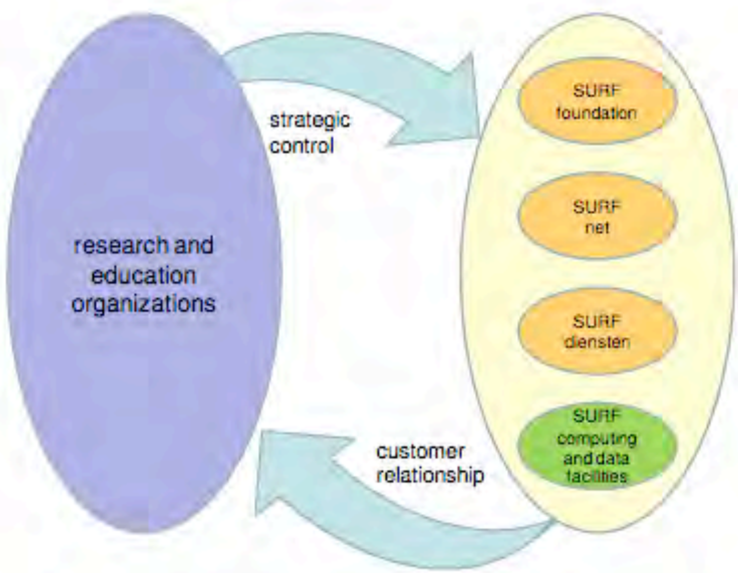

Figure 3 Proposed SURF organization

In SURF a new entity is envisaged. This entity could be named Computing and Data Facilities (in Dutch *Rekenfaciliteiten*, referring to the original meaning of the name of SURF...). This entity is responsible for the development, operations and support of the e-science services, grids, and resources (computing resources, data

To obtain the entire report:
http://www.ictregie.nl/publicaties/nl_08-NROI-258_Advies_ICT_infrastructuur_vdef.pdf

Executive Summary

Introduction pp. 1-3

This issue examines the continued evolution of both fiber infrastructure and optical networking in research, education, and tech transfer networks and network projects in the Netherlands. It also looks at the Dutch planning process for directing economic investment in ICT in ways designed to achieve, through innovation, maximum economic impact.

The introduction and conclusion extensively discusses the programs of ICTRegie – the Dutch ICT Research and Innovation Authority. They reflect a process set up in 2005 whereby Government research funding is put through a process that is designed to create proposals that are in keeping with nationally defined goals allied to the nation's social needs as well as its economic strengths.

The mission statement is what President Obama should emulate in his Office of the CTO. "First to introduce unity and consistency to the strategic direction of the ICT research and innovation by means of the development of a national strategy that enjoys broad support and, second, to ensure ongoing strengthening and appropriate dynamism of a Dutch ICT knowledge infrastructure geared towards high social and economic yield."

A New National Infrastructure, pp. 4-14

Kees Neggers, SURFnet Managing Director, explains how innovative technologists used the intersection of smart Dutch telecom liberalization policies, and then the dot com boom to build competing fiber rings around Amsterdam creating an attractive environment that allowed the city to become a global hub for layer three traffic at AMS-IX and layer one and two optical light waves at NetherLight.

Kees explains the sequence of events leading to SURFnet's solving the last mile connection by means of fiber for all its members in 2004 and becoming the first national all optical research network at the end of 2005. This enabled what they call hybrid networking – doing as many tasks a possible at layer one and two and as few tasks as possible with much more expensive electrical and routed layer three networks.

Working with an all optical network made the assignment of light paths possible, first to institutions, then to departments and ultimately to individuals.

The five year GigaPort project running from January 1 2004 through December 31 2008 made optical private networks possible for SURFnet members. It also made possible the establishment of a resource management platform built on a foundation of grid computing, a layer of web services allowing users to connect to applications and to instruments as well as super computers. On top of the resource management platform is one of generic e-science services (the light blue box in the diagram on page 11) in other words software tools that support research in individual fields.

Within GigaPort one has cutting edge technology supporting not only basic research but also tech-transfer from universities to companies like Philips, Unilever and IBM. It also has valuable experience with the complexities of making it work through ways to build and administer the corporate networks, to the software tools needed to take advantage of the lightpaths to needed training to enable users to benefit from the technology at their disposal.

Most significant is the decision of ICTRegie to recommend continued support for SURFnet and GigaPort as a basic platform for technology development in the Netherlands. As Harvey Newman says at the very end of this issue, while the private sector is critical without roads open to all we will be placing a huge handicap on our possibility for further progress.

As Kees Neggers said: **In our view innovation needs the network and hardware resources plus the software and the people that have the knowledge of how all of these layers interact with each other and in addition to this knowledge you still need a major outreach effort to involve the users for which it has been created.**

[And Cees de Laat added]: **From the light blue box you need people working with every discipline to make it happen, because the people in the disciplines themselves cannot do that. You can not expect the biology professor to understand how to go out and identify and hire the people necessary to teach him and his colleagues how to use the high performance tools.**

Hybrid Networks and e-Science Future Development, pp. 15 -40

Cees de Laat takes us through the events that began with SURFnet's 2001 decision to work with Starlight in Chicago and CENIC in San Diego as well as a few other locations to explore the use of dedicated light waves of one to ten gigabits. The experiences of the iGrid2002 meeting led Cees and his colleagues to begin work with circuit switched layer 2 lightpaths wherever possible.

Working with extremely demanding users who needed their own lightpaths led Cees and his group to think about tailoring the network to the needs of the user so that, unlike layer 3 routed networks that are stupid procrustean beds into which everything must fit, the networks could become much more flexible and adaptive for the population that the network is designed to serve.

As Cees explains: "what I am describing is a kind of peculiar programming language or programming environment where your networks are just subroutines, your data are more subroutines, your solutions to solve the problem are also subroutines. You then say optimize and solve and it will then work out the most optimal way to achieve what you want to have done.

What all this means – and this is the most fundamental thing to get one's mind around – is that your network becomes just part of your programming environment.

Normally you have your data and your computing and a "stupid network" that is an unmodifiable "given" that you have to play along with. **But here your network is just part of your toolset."**

Later he elaborates: "In our group back home we are creating objects for virtualizing and programming wavelength switches and photonic devices which directly talk to fiber; Micro-Electro-Mechanical Systems (MEMS) devices that can connect fibers so that we have flexibility at the fiber layer. And at the Ethernet layer, we can do similar things manipulating Virtual Local Area Networks (VLANs). We address the Ethernet layer and the packet routing when we need that.

If you can manipulate all these layers and have also vertical and horizontal knowledge in every layer, you can do the magic and you get a perfectly integrated multilayer hybrid network that is optimized for your application."

Cees goes on to elaborate the kinds of e-Science projects that these lightpath networks support. He then demonstrates in some detail his VL-e (e-Science Virtual Laboratory) applications.

I said to him: what you are explaining is **how to build a networking environment that combines this use of equipment, technology, and energy in such a way to give the user applications that meet his demands at the minimal necessary cost?**

De Laat: "Yes. But to do this one must expose the network to the users so that they understand what is happening and can see the benefits of not always working at the highest and most expensive layer -- namely the routed layer. One needs to make sure that your users understand the benefits of operating in different layers of their optical network. The users need to grasp that if they endow their applications with the ability to intelligently traverse the layers they can open up all manner of increased possibilities to better performance at less cost."

For the most part the rest of interview describes how routing or path finding is done in his hybrid optical network.

Interactive Networks, p. 41-49

I have written up with illustrative photos a demonstration by Rudolf Strijkers of what a user interface to application embedded networks looks like and the capabilities it permits.

What we discuss is a prototype of a user interface with a multi-touch-sensitive screen that allows a user to tap on the tools he wishes to select and with his finger to draw the paths he wishes in order to activate nodes on a programmable network.

As Rudolf pointed out: This is the first prototype of what we are calling "Interactive Networks". In interactive networks humans become an integral part of the control system to manage the next-generation of programmable networks and Grids. The main design principle is this; by virtualizing the configurable and programmable properties of network elements as software objects, any aspect of a network infrastructure can be manipulated from computer programs. What we show here is an implementation of an interactive control system concept for user programmable networks, which applies the architectural concepts we have developed in our research.

At the end he concludes: **With the information given by the network, and currently we support continuous measurement of delay, jitter, bandwidth and throughput, one or more operators at once can write programs that automate decision-making. This opens the way for automated network adaptation in a user-friendly environment.**

For example, now we can say, certain paths should avoid busy parts of the network. By using only standard functions in Mathematica, it is already possible to write a simple program that uses the real-time throughput information to continuously reroute one or more paths that avoid busy parts of the network.

Conclusion pp. 50-58

The technology that is discussed here is inevitable and inexorably extensible and usable at reasonable cost where open access fiber is available. The countries that implement this openly and broadly as an infrastructural system of roads and highways will derive enormous competitive educational and economic benefit. Those that don't will never see either the educational or the economic benefits.

We touch on the OptiPuter and OptiPortals and their significance for science education. We also extensively discuss the Netherlands ICTRegie five year plan and show how it can be used to further distributed edge based collaboration and consensus in allocating economic resources. We offer a seven point national fiber infrastructure plan for the new administration.

Be it a new ARPA or an American ICTRegie, our new administration MUST grasp the present opportunity to sweep out the wreckage of cronyism and speculative capital. In the national and public interest like Roosevelt did in the 1930s, it must build out a "rural electrification system" - this time an interstate highway system of optical fiber.

I conclude with a mention of BT becoming a services based cloud computing platform that I made to Harvey Newman.

Harvey Newman: I agree with this. And it makes me think - Often, to make one's case for innovating at an incumbent, one is constrained to say "What is the business model of the Internet ?"

In this case what they are asking is: How do companies ("how do I") make a profit out of this ?

So I say - Only the highway-building model ("building the nation from the bottom up") will get us out of the past.

And if we do not do this - there is strategic as well as economic risk. Since any nation can do this."

COOK Report: I will be further exploring the implications of this via an interview with Harvey in either the March or April issue.

A Note from the Editor on the February 2009 Format and Presentation

This issue has an Introduction; an interview with Kees Neggers; another with Cees de Laat; a presentation by Rudolf Strijkers and a closing essay. The Symposium discussion is postponed for a month

Text, URLs and Executive Summary: I have attempted to identify especially noteworthy text by means of boldface for **REALLY** good "stuff" . **Also the proper Executive Summary in this issue continues**. I hope you find it useful. Feedback welcomed. You will also find live URL links and page links in this issue.. (I am also no longer changing British spellings of things like fibre to the American fiber.)

Thanks to **Sara Wedeman -** see sarasworld.blogspot.com/behavioraleconomics/ for assistance with the masthead logo. Captain Cook now charts direction by looking at a compass rosette.

Coming in the March 2009 issue - out about January 31 probably a discussion with Frank Coluccio on fiber to the desk top and taking copper out of the networks. On December 29th I completed a marathon 4 plus hour interview with Harvey Newman. The March issue will definitely have symposium discussion - I am not sure at this point how i will handle the Newman and Coluccio material

I am omitting the contributors' page since a cumulative list may now be found at http://www.cookreport.com/index.php?option=com_content&view=article&id=121&Itemid=74

Contents