BUILDING A NATIONAL KNOWLEDGE INFRASTRUCTURE

HOW DUTCH PRAGMATISM

NURTURES A 21ST CENTURY ECONOMY

*The COOK Report
On Internet Protocol*

Building a National Knowledge Infrastructure

How Dutch Pragmatism Nurtures
a 21st Century Economy

Gordon Cook

The COOK Report on Internet Protocol

Dedication

To Valerie, Gordon III, Annalee, and Max my family
whose support made this possible

and to Carlota Perez whose magnificent work provided
the missing link and thereby enabled the connection that
provides the foundation for this book

with warm appreciation

Gordon Cook

The Origin of Knowledge Infrastructure

The internet and broadband debate in the US is too close to the "ground" to enable most observers to understand where the world will be in five years. That reality five years off is the knowledge infrastructure being built in the Netherlands and described in this book. This far-sighted work will be transformative there in the same way that water power and steam power transformed England in the 18th century. This while the rest of the world is arguing over the private share-owner model for what should be public infrastructure. This debate deprives it of the most important advantage of optical technology. When it makes this transition it will be able to begin where the Dutch started in 1997 with the parliamentary motion favoring the building of knowledge infrastructure. Meanwhile the Netherlands has a 13 year lead. As long as the Dutch keep gong the rest of the world will never catch up.

Nations that don't understand this will condemn their citizens to be the poorly paid peasants of the underdeveloped world of the 21st century. If the United States doesn't learn these lessons it will be transformed into the world's leading digital back water. Rather like the contrast of 19th century Russia. The intellect and brilliance and economy was all in either Moscow or St Petersburg. Elsewhere was the largest landmass of backward poverty on the planet. The US is still blessed. I am fortunate enough to have some of the better broadband – without this I can say from first hand experience I could never have accepted the assignment from the Netherlands that turned a double issue of my newsletter into a book being able to communicate with Europe instantaneously 3000 miles and six time zones to the East and with my Editor Sheldon Renan in Oregon 2000 miles and three more times zones to the West in the space of 10 days.

I have been reporting on SURFnet since Bill St. Arnaud introduced me to Kees Neggers at a CANARIE Workshop in Montreal in 1999. My question to Bill was who is doing this right? Who has a philosophy similar to yours? His response – "Kees Neggers" and the in-person introduction was accomplished. I have been reporting on SURFnet ever since. At the end of 2004 and beginning of 2005 I did a long report. In November 2008 I met Kees again at Supercomputing 08 and did a much longer report with him and Cees de Laat because I was finally gaining an appreciation of the very large importance of what they are doing. A year later at Supercomputing 09 we met again, this time also with Bob Hertzberger, Hans Dijkman and Wim Liebrand.

I found the information necessary to answer the questions that came to mind and wove them into the fabric that has become this book. A fabric that was initially intended only for issues of my newsletter.

As a matter of process I fact–checked what I was finding with the folk I interviewed in Portland. They answered my questions, liked my drafts and wanted them to be read by more than just my subscribers and hence asked me to collaborate in turning them into the resulting book.

This book's intent is to lay out a roadmap that explains to readers how the Netherlands came to its national consensus to invest in national knowledge infrastructure. And show what the evolution of the resulting policy process has accomplished. It will conclude by pointing out the changes in perception that must be accomplished in the united states in order to render feasible what Holland has done. Namely treat the fiber and ducts as the national highway system. Those are held in custody for the use of the nation and managed by the government which services the citizens for the benefit of the people and on which private enterprise can build competing higher layer, profit making operations.

Only then can any nation build a competitive, economically valuable knowledge infrastructure. The can become possible in the United States only if competing interests simply understand the policy principle of infrastructure as laid out in the chapters of this book.

Gordon S. Cook, Jr.
Ewing, NJ USA
February, 3 2010

Acknowledgments

To **Valerie** my wife of 43 years for her good natured toleration and love for whom I have become and for being a woman of grace and compassion.

To my son **Gordon III** whom I have watched become a hugely talented builder and shaper of environments who has been there in ways beyond comprehension.

To my daughter **Annalee** who has become a warm independent strong loving mother for my grandson "Max". And for Max who must inherit the digital infrastructure described here in order to thrive.

To the **Colonel, Dave Hughes**, West Point Class of 1950 also known as the "cursor cowboy' who has been a good friend since 1980 and who when I finished my policy work at the US Congress Office of Technology Assessment on March 1, 1992 when I asked what the hell will I do now? said without missing a beat - write and publish a newsletter about this internet thing. Dave saw the synergy possible with convergence of my skills and the technology. Standing on the cusp of the 19th year of the *COOK Report* I have never before been so excited about the possibilities.

To **Sheldon Renan** whom I have known now for almost four years.   Man of renaissance talents, writer, producer, designer, poet, master story teller and seer. Finder of the rare essence in our sea of chaos. The spirit who has guided me to connect a carpet of events and ideas that want to be related want to be connected and catalysed them into the

following tale. Just as he has been weaving in his own mind for two decades the concept of 'netness' or Renan's law. - Everything wants to be connected and works better when connected. This book is a specific example of the deeper drivers of Renan's Law.

To **Sara Wedeman** one of the most creative, determined, talented women I have ever met who lives in the trans-disciplinary mind set that the world may get to five years from now. Psychologist, econometrician, down loader of the complete databases of the US Census, and practitioner extraordinaire of behavioral economics. Visionary proponent of broadband mapping and Amartya Sen who, if ten years from now there still is an NTIA, Sara will lead the agency to act in the interest of our nation. And an artist who generously devoted more than four full days of design work to get the cover of this book done right.

and

To **Eva Waskell** who in the Spring of 1992 discovered that there was someone at OTA trying to describe this Al Gore research network and invited me to talk at a Tele-strategies conference in April 1992. And whom, as I only found many years later, 25 years ago began what has morphed into the nation wide election integrity movement. No matter what the obstacles Eva has never budged.

All visionaries, all determined, all motivated by the public good. All Bodhisattvas.

Contents

COOK REPORT
ON INTERNET PROTOCOL
FEBRUARY/MARCH 2010

I. How the Dutch Did It Right

When so many others get it wrong

In 1965, French film director Jean-Luc Godard made a film noir called Alphaville about a grim future in which all the world's infrastructure is controlled by a computer. Imagine that. Towards the end of Alphaville, the hero pulls the computer's plug. Suddenly nothing works. The world goes blind. People are feeling their way along walls.

In 1965 infrastructure controlled by computer was science fiction. And today?

Today lives and economies depend on the Information and Communications Technology (ICT) Infrastructure more every day. The quality of the national ICT infrastructure increasingly defines the quality of life as well as opportunity for the future.

It used to be the United States, with it's Internet initiative and "intelligent networks" that claimed the leadership position. Now that's the science fiction.

Today it is the Netherlands which has consistently taken the initiative, made the investments and delivered the goods. The Dutch have developed blazingly fast hybrid optical networks of hitherto unimagined efficiency and cost effectiveness.

They have taken the technique of user controlled light paths as developed by CANARIE in Canada and incorporated this on a robust platform of Web services and with collaborative e-science middleware into optical network technology. So their networks cannot only transmit huge data files across oceans starting one side of the world to the other instantaneously. But increasingly the Dutch can support the collaborative and multidisciplinary practices that are being described as "the 4th paradigm", data intensive science.

Unlike the United States — where private interests have walled off and Balkanized

much of the Internet — the Dutch are committed to collaboration both inside and outside the country. They have been a proactive force for international collaboration with scientists and network specialists in the EU, the United States and elsewhere with their Global Lambda Integrated Facility or GLIF.

In 1965 infrastructure controlled by computer was science fiction.

And the more you learn, the more it is apparent that the Netherlands is building the electronic network and knowledge infrastructure on which the economy of the 21st century will be based. This report will explain their continuing technology direc-

tion -- a direction that on its own level is quite impressive.

So why is this happening in a country of only 17 million people, small and crowded with one of the highest population densities in the world? Why is this not happening in the United States or Canada or China or Brazil? Those are all big countries, rich in resources, that are supposed to "own the future." This is an important question.

We can begin with what anybody who has worked in technology knows. The real problems with technology are not, typically, the technology. The real problems have to do with intentions, models, governance, implementation, resourcing, and user support. It's not the technology, it's what you do with it. And that, in turn, depends upon what you have intended to do with the technology.

So when we ask the question, why are the Dutch doing it right, when so many other countries got it wrong, the answer begins with the fact that they have better intentions and a more inclusive process. In demonstrating admirably farsighted planning and negotiated discussion among their stakeholders, the Dutch are leading the world in making the ICT technology transition.

This is the change described by Carlota Perez that is common to all technology revolutions. Speculative or finance capital in the support and development of ICT must no longer predominate. Society must shift to use of productive capital. In other words it must use money for infrastructure to install these ICT resources in society and treat them as knowledge infrastructure. They become the logical follow on to roads and highways, canals, railroads, electric grids, airports water and sewage systems and electric plants. In the end they are simply an integral part of the basic infrastructure of an advanced and civilized capitalist nation.

…the Netherlands is building the electronic network knowledge infrastructure on which future economies will be based.

So, once again, it is not the technology so much as it is the thoughtful and careful way that technology policy is determined. It is the way that policy is turned into reality. And it is the way in which they continue to push edges

in a never ending pursuit of technical excellence and extreme performance. And it is the counter-intuitive strategy in which geek values of open source software and collaborative networks are harvested into creating public/ private partnerships that evolve into the business opportunities and a dynamic and competitive national economy.

…embrace disruption and use the improved productivity for the public good.

John Hagel and John Seely Brown recently updated their sobering 2009 Shift Index: Measuring the forces of long-term change, which reveals that in the U.S., despite an economic focus on private good rather than public good, American businesses, including telcos, now earn 75% less return on assets than they did in 1965. One of the reasons, according to John Hagel, speaking at the 2009 SuperNova Conference in San Francisco, is that businesses have not been able to rationalize the disruptive advances in technology which (thank you, Moore's Law) never stop advancing.

The severity of events is made more difficult because of the doctrine that the independence of the private carrier is sacrosanct. The problem is that the direction of technology has been moving ever more rapidly into the creation of network capability of almost limitless abundance. In contrast to this technology push, the investment pull of the rules by which the share owner Corporation is governed demands that management take the opposite course and establish a regime based on scarcity - measured usage, constricted bandwidth, constricted user freedom and charging the user the very maximum that traffic will bear. The privatized carrier becomes a predator that feeds on society that in theory it serves. The result is an enormous gap between the technological capabilities attainable with state-of-the-art optical network technology and the reality enforced by share owner networks in their respective societies

In the Netherlands, the ICT infrastructure ecosystem does have a strategy. It works for them. And it works for us. And that is to embrace those disruptive advances and use the improved productivity for the public good.

The COOK Report contends that the Netherlands, through the good fortune of rather unique circumstances over the past decade, has been able to articulate a vision whereby it begins to treat its information and communication technology investment as an investment in public infrastructure rather than private share owner determined enterprise. The nation telecommunications infrastructure is treated as a public rather than a private good. Partly this is for the purposes of science and research. But it has allowed the country to develop world leading technology that operates alongside the share owner maintained voice network of KPN as well as those of the MSOs (Cable TV companies) of the Netherlands.

In trying to understand the emergence of the Netherlands as a leader in ICT infrastructure, again, this is not just about network technology or the network applications. Rather the key lesson to take away is that the Dutch network and research effort has been built by means of an exceptional attention paid to the economic impact of the network as infrastructure that contributes to the Dutch national interest.

And beyond that is another parallel story. It's the story of how history and economic circumstance shaped the character and confidence of the Dutch. It tells what can happen when a nation, toughened by centuries of challenges that would (literally) sink most countries, learn how to work together to leverage technology at the infrastructure level and overcome impossible odds. It's a great story, and it starts in Chapter II.

A Question for the United States: How to Make Benefits of Thinking Long-term More Visible?

The long-term vision and achievements of SURFnet and the Dutch infrastructure ecosystem call into question the failure of U.S. infrastructure players to be equally committed to the long view. One question we need to ask, as a nation, is this. How do we make long-term risk and benefits more visible so that people and institutions can make better choices?

Failure to understand true long term risks and opportunities is responsible for many investment cycle failures. Most recently the mortgage crisis, but more commonly the cycle in which insurance companies lower their rates in lock step to stay in business, only to find during some catastrophe that their underwriting models were too optimistic. On investigation, they invariably find that (a) the models really weren't very good, but (b) the risk assessment assumptions in the models had been progressively downgraded in response to competitive pressure. This is not a consequence of stupidity. It is a requirement for market survival. Or to build a great infrastructure.

Johnathon S Schapiro wrote to the IP list on Jan 4: "While bank lending practices were, in my opinion, largely responsible for the mortgage crisis, it must be acknowledged that banks, trading houses, and insurers suffer in common under what might be termed 'the competitive death embrace' ".

The death embrace is best illustrated by the thinking "If I don't do this marginal deal, my competitor will, so I should do the deal rather than let the benefit go to them." If you review the papers, you'll see countless variants of that statement, and if you pay attention, you'll notice that not one says "benefits and risk".

This problem is compounded by the short-term biases of corporate securities reporting. A Warren Buffet understands this very well and reads through it. (Many readers of the Cook Report understand this but most lack the investment discipline to profit from it.)

In technical contexts, we see companies and decision -makers repeatedly caught in what Harvard's Clayton Christensen has termed the "Innovator's Dilemma."

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

Disruption, whether in markets or technologies, is inevitable. Too many choose to ignore that fact. The cause can ultimately be traced to the fact that quarterly behavior can be explained to investors while long-term behavior can't.

II. The Netherlands National ICT Research Infrastructure

History, character, policy and pragmatism

Infrastructure and the National Interest

After World War Two, government in both the United States and Europe continued the social contract that had helped end the problems of the Depression. In Europe, however, it was necessary to rebuild infrastructure from the ground up. The flourishing of free market capitalism was kept in balance by a generally accepted consensus that there needed to be a social safety net, and that government needed to create the conditions under which its citizens had basic food, shelter and education.

The Dutch have done an extraordinary job of determining how to leverage their resources in order to create a knowledge infrastructure that will serve the national interest, support EU partnerships, and compete successfully in the international arena. Meanwhile in the United States, infrastructure discussions and investments have been sporadic. The U.S. has not pursued a serious infrastructural agenda since the interstate highway system was constructed in the 1950s and 60s.

… after the war infrastructure had to be rebuilt from the ground up.

The Economy of the Netherlands

Wikipedia says "On the Index of Economic Freedom, the Netherlands is the 13th most laissez-faire capitalist economy out of 157 surveyed countries. At the time of writing the Netherlands is the 16th largest economy of the world. Between 1998 and 2000 annual economic growth (GDP) averaged nearly 4%, well above the European average. Growth slowed considerably in 2001-05 as part of the global economic slowdown. 2006 however, showed a promising 2.9% growth. Yearly growth accelerated to 4.2% in the third quarter of 2007. Inflation is 1.3% and is expected to stay low at about 1.5% in the coming years. The Netherlands has a prosperous and open economy, which depends heavily on foreign trade. The economy is noted for stable industrial relations, fairly low unemployment and inflation, a sizable current account surplus, and an important role as a European transportation hub. Industrial activity is predominantly in food processing, chemicals, petroleum refining, and electrical machinery." http://en.wikipedia.org/wiki/Economy_of_the_Netherlands

Role of Natural Gas in Dutch Postwar Good Fortune

The Netherlands has, however, an asset that for a limited time continues to generate money for investing in social capital. The Netherlands is presently the world's fifth-greatest natural gas exporter. According to Wikipedia, "While its oil reserves in the North Sea are of little importance, the Netherlands is presently the second-greatest [after Norway] natural gas producer in the European Union and the ninth-greatest in the world, [after Russia] accounting for more than 30% of EU total annual gas production and about 2.7% of the annual world total. Proven natural gas reserves of the Netherlands are estimated (as of January 2005)

at about 50-60 trillion cubic feet, or about 0.9% of the world total. Although the Netherlands owns substantial gas reserves in the North Sea, most of its production is presently from on-shore wells, and much of the natural gas produced by the Netherlands comes from Groningen Province, which borders the North Sea. [Same Wikipedia article cited above.]

Drilling began in the Slochteren natural gas field - one of the biggest in the world on July 22, 1959. Fifty years later according to a June 17 article in NRC Handelsblad "the total revenue from the Groningen gas field is more than 211 billion euros, most of which went straight to the treasury." http://www.nrc.nl/internation al/article2274261.ece/The_D utch_curse_how_billions_fro m_natural_gas_went_up_in_ smoke
"The Slochteren natural gas field also turned out to be much bigger than expected. Back in 1963, the NAM estimated reserves at 1,100 billion cubic meters. That number has been revised upwards to 2,700 billion cubic metres, more than 60 percent of which has been extracted. . . . "more than 52 billion Euros, or almost a quarter of all natural gas revenues, went into financing social security. Only 15 percent was used to improve the

national infrastructure of the Netherlands, while 85 percent went to welfare benefits, interest payments on the national debt, and spending on health care, education and public administration."

"Ruud Lubbers, who was economic affairs minister in the seventies and prime minister from 1982-1994, complained that the money led to a "lack of discipline".   It wasn't until the late eighties that Lubbers saw an opportunity to create a fund for structural improvements to the economy. The "gas fund" (Fonds Economische Structuurversterking, FES) finally saw the light in 1993."

In many countries, revenues from natural resources invariably are siphoned off by private interests. The Dutch have had the discipline to turn the income from gas to infrastructure and then to knowledge infrastructure. The difference may be the innate pragmatism of the Dutch character. For centuries they have faced unending challenges that have required working together and managing a commons involving water resources and innovative technology. Everyone knows the story of the dikes and windmills. Fewer know of the breakthrough invention of a windmill-powered sawmill that enabled the Dutch to build a flotilla of smaller, faster, and cheaper ships that

then defeated the mighty Spanish armada.
http://en.wikipedia.org/wiki/ Cornelis_Corneliszoon
The COOK Report contends it is this heritage that has enabled the Dutch to begin their critical and innovative use of the gas funds in building the knowledge infrastructure that the remainder of this chapter will describe.

On the Other Hand the United States Plies a Different Course

In the United States, the Rooseveltian New Deal put in a framework of financial regulation that enabled the U.S. to prosper after World War Two. In the 50s, under President Dwight David Eisenhower, national prosperity increased during a period when the top income tax rate was what today would be an unthinkable rate of 93%. Under President Kennedy, however, that rate was cut to roughly 70%, The course of the next four decades was increasingly set by the increasing triumph of the Chicago School of free market economics which preached that economic growth should be achieved by more and more tax cuts. This resulted in increasing disparity in financial reward for the work force. CEOs salaries grew from a multiple of 20 times entry level worker earnings to a multiple of several hun-

dred times entry level compensation. In 1968 the CEO of AT&T made $200,000. Today chief executives at AT&T and Verizon receive salaries of up to $30 million. After bonuses, their annual incomes may exceed $100 million dollars per year.

…challenges required the Dutch to work together to manage a commons …with innovative technology.

The dominant ideology in the US became that the government was "bad" and needed to be starved. On the other hand, success of private industry and the so-called free market "would raise all boats". In reality it meant that the American government had limited ability to act in the national or public interest. When given a choice between funding infrastructure or funding the military, military always won. This trend reached critical mass during Ronald Reagan's presidency. It enshrined the supremacy of private interest over public good. This resulted in significant increases in the national debt as well as the financialization of the economy. Growth was fueled primarily by credit, while productive activities like

manufacturing hugely declined. The manipulation of money made up an unprecedented share of U.S. GDP.

Ironically this phenomena is sometimes described as "the Dutch disease." The term was coined by The Economist in 1977 to describe how the discovery of natural gas in the Netherlands resulted in a significant decline in manufacturing there. That said, the Netherlands learned how to manage it. The United States did not.
http://en.wikipedia.org/wiki/Dutch_disease

…in the U.S. private interest was put before public good.

As the integrated circuit (invented in 1958) matured and made possible the global internet, speculative capital funded that growth in the United States, giving it significant control of the national economic agenda. The result was the extension of what Carlota Perez in her seminal 2003 study of technology revolutions defines as "the transition phase" needed to shift national policy from necessary investment in producing a new technology to the more efficient use by national government of productive capital.   In this state,

Perez points out, economic policy decisions made in the national interest, as opposed to supporting Wall Street, can result in disseminating technology and its benefits across the broadest reaches of society.

See:
http://en.wikipedia.org/wiki/Carlota_Perez
and
http://www.carlotaperez.org/Articulos/TRFC-TOCeng.htm
and
http://www.cookreport.com/carlota.pdf

But this is not what happened in the U.S. and Britain. Reagan was elected in the U.S. at the same time as Thatcherism triumphed in the United Kingdom. The incumbent carriers were privatized along with the early cable tv systems where growth in the EU was far behind the US. With support of incumbent telcos, economic market conditions were established to obstruct effective use of new optical fiber technologies.

The concept of a general purpose technology (GPT) is useful in understanding this issue.   GPTs can affect an entire economy (usually at a national or global level) and have the potential to drastically alter societies through their ability to disrupt pre-existing economic and social structures.

...incumbents obstruct effective use of optical fibre technology

Mark Cooper, Director of Research for the Consumer of Federation of America wrote on January 8 2010 about the potential benefits of dissemination GPTs in the United States. "The extremely diffuse benefits mean that the dominant funding mechanism of the past thirty years (private capital) won't work very well. The approach to funding the last three GPTs in the U.S. was very mixed and pragmatic. Railroads were first funded in the private sector and they all went bankrupt. We then funded them with land grants, lucra-

tive mail hauling franchises, etc. Electricity and telephone were done primarily as regulated franchises (until the new deal when it was realized that a large part of the nation would never be served so we set up cooperative and municipal entities). Free market ideology still dominates U.S. policymaking. Until we have overthrown that ideology, there is no way to fund a ubiquitous, adequate national broadband infrastructure."

The SURF Foundation

As privatization of KPN, the Dutch incumbent, approached, and small urban cable systems were sold off, far-sighted people in the Netherlands were quick to appreciate the development

of early packet data networks and the need for having a coordinated government approach to the creation of a university and research-oriented infrastructure. As a result the SURF foundation was created along side the NWO or Dutch Science Foundation.

"The SURF Foundation was established in 1987 to coordinate the implementation of a multi-year plan for the improvement of the application of information technology (IT) in Dutch universities, schools for higher vocational education and research institutes. In the course of its activities SURF has become a nationwide supplier of services. These services are primarily provided through its operating subsidiaries: SURFnet bv and SURFdiensten bv. SURFnet manages the computer network of the same name. SURFdiensten deals with licensing agreements in the fields of software, hardware and information services."
http://www.ariadne.ac.uk/issue5/surf/

The SURF site itself says: "SURF is the collaborative organization for higher education institutions and research institutes aimed at breakthrough innovations in ICT. SURF provides the foundation for the excellence of higher education and research in the Netherlands."

"On 13 June 1986, RARE (Réseaux Associés pour la Recherche Européenne) **is established as an association under Dutch law by Hans Rosenberg on behalf of the University of Utrecht and Klaus Ullmann on behalf of the DFN Association. This photo and caption is found in"Terena Celebrating 20 years" booklet: http://www.terena.org/publications/files/20th_anniversary.pdf**

http://www.surffoundation.nl/en/oversurf/Pages/Default.aspx

Kees Neggers, SURFnet Director says that "SURF existed already [before 1987] as an informal approach to cooperation between the Universities with regular meetings in which the University Board member responsible for ICT and the Director of the Computing Centre participated." It was this group that took the initiative in the mid eighties to start a national research network as part of a broader ICT plan for the research and higher education community.

...Hans Rosenberg understood why a small country needed a big network.

The Government reacted positively to this informal initiative by first financing the start of SURFnet as a project and two years later, under the condition that SURF would first become a formal legal entity to create SURFnet as a company in cooperation with the PTT (at that time the state owned telecom monopoly operator and since 1989 privatised into KPN), the government also funded the start of SURFnet as a permanent organization."

Leading the initiative was the visionary Hans Rosenberg, he was member of the board of the University of Utrecht. Before that he was an alderman in Utrecht and was prominent in solar radio astronomy. Rosenberg was the founding father of the current ICT infrastructure in the Netherlands. He helped start SURF,

The Development of AMS-IX

The Science Park Amsterdam (the relatively new name of the location of the networking institutes and AMS-IX and Netherlight) in is a historical networking hub. By 1982 CWI the National Research institute for Mathmatics and Computer Science had become a European hub in EUnet with its famous "mcvax." In the early nineties, NIKHEF, with Rob Blokzijl as its driving force, and SURFnet at SARA formed a layer-2 shared infrastructure to exchange traffic between (academic) organizations. In February 1994, it was internationalized as a community based Internet Exchange Point to exchange traffic with CERN in Switzerland.   At this point other ISP's were also allowed to connect and the name AMS-IX was first used. In 1997 NIKHEF and SURFnet handed over the formal responsibility of the distributed exchange to the AMS-IX Association. In 1990 the Association formed the AMS-IX limited company, AMS-IX B.V.   The Association held all shares while all assets are transferred to the company AMS-IX B.V. Until 2002 SURFnet continued to manage the overall operations of the exchange.

Today AMS-IX still operates as a not-for-profit community based exchange with well over 300 connected members. With a daily average traffic over 500 Gbit/s and peak traffic over 800 Gbit/s AMS-IX today is the largest internet exchange in the world.
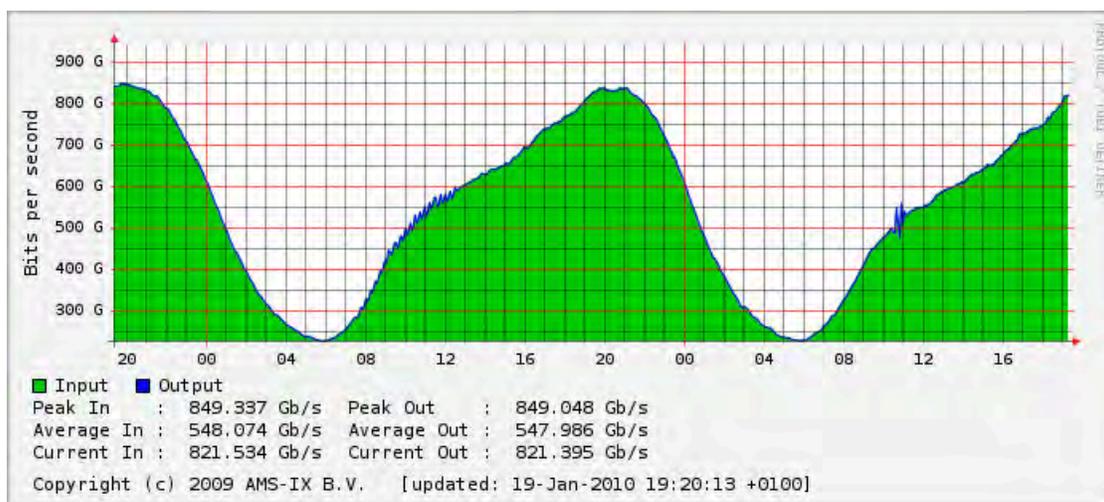
NetherLight: a Layer 1 - 2 Optical Exchange

NetherLight is a SURFnet owned and operated advanced open optical switching infrastructure and proving ground for network services optimized for high-performance applications. NetherLight has been operational since January 2002 with the installation of a 2.5 Gbit/s trans-Atlantic lambda between NetherLight and Star-Light in Chicago. NetherLight is a SONET/SDH cross connect and Gigabit Ethernet switching facility for high-performance access to participating networks. Ultimately, it will become a pure wavelength switching facility for wavelength circuits as optical technologies and their control planes mature. NetherLight is made possible by SURFnet in the context of the GigaPort project. NetherLight is located at SARA in the Science Park Amsterdam, and has been realized by SURFnet in the context of the GigaPort projects. Today over 300 Gbit/s of international lightpath capacity is connected to NetherLight. http://www.netherlight.net/

SARA

SARA Computing and Networking Services is an advanced ICT service center that supplies – since more than 30 years – a complete package of high performance computing & visualization, high performance networking and infrastructure services. Among SARA's customers are the business community and scientific, educational, and government institutions. http://www.sara.nl/aboutsara/aboutsara_pr_eng.html

SARA's product portfolio consists of:
- High Performance Computing & Visualization: facilities and services in supercomputing, data storage, visualization and virtual reality;
- High Performance Networking: design, installation and management of advanced Wide Area Networks;
- ICT Services: housing and management for third parties of critical infrastructures, systems, applications
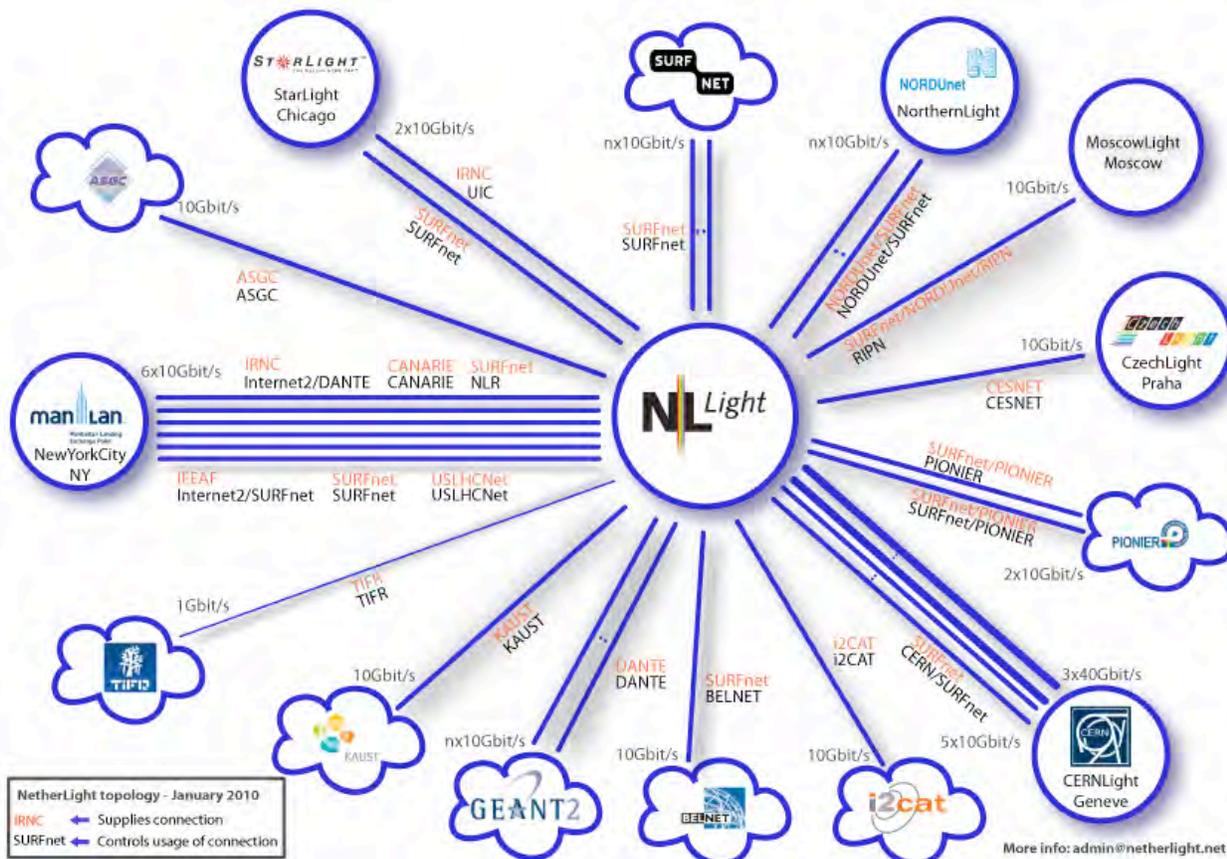
The above diagram shows the January 2010 international connectivity at NetherLight, the GLIF Open Lightpath Exchange (GOLE) in Amsterdam. All other circles in the diagram are also GOLEs like NetherLight, the clouds are the networks connected to NetherLight that are capable of supplying lightpaths. Details of all GOLEs can be found on the GLIF website under Resources: http://www.glif.is/resources/

SURFnet, and was also Founder of RARE.   And it was Hans Rosenberg who understood why a small country needed a big network. Rosenberg died untimely in 1992, but his ideas and influence are still felt by the current leadership of SURF

SURFnet started as a project in 1985/1986 Stichting SURF was created in March 1987. In January 1988 SURFnet became a not-for-profit company 'in formation' (which means SURFnet operated as a company, but was not yet a formal legal entity, SURF was formally responsible for its business.)

In January 1989 SURFnet became a formal legal entity, SURFnet bv, with SURF having 51% of the shares and KPN 49%. (One of the reasons for the delay was that it was only in January, 1989, that PTT became KPN as a private company. Before that it would have been more complicated for the PTT to take shares in a company).

In 1999 GigaSURF bv was started, 100% owned by SURF, to do the request for tenders of the SURFnet5 network as part of the first GigaPort project, SURFnet continued to do the actual work. (This construction was needed because KPN initially was unwilling to give up its shares after the liberalization of the Telecoms market was

completed, and it was considered not correct to have KPN reacting to request for tenders from a company in which it owned 49% of the shares).

On 31 December 2001 KPN sold its shares in SURFnet to SURF and the next day SURFnet and GigaSURF merged and GigaSURF was abandoned. Since then, SURFnet bv is a 100% subsidiary of SURF."

...SURFnet converts natural gas into ICT infrastructure

Transition to and Growing up with Infrastructure

SURFnet is important for several reasons. The most important are its consistent services and persistent visibibilty – and especially its role in enabling the funding of the ICT infrastructure via natural gas revenues in the 1990s.   That SURFnet and related entities like AMS-IX and NetherLight are just "there" now and taken for granted is perhaps the single largest accomplishment that the SURF Foundation has made during the last 15 years.
http://wiki.glif.is/index.php/NetherLight

On January 5, 2010, the 200 members of the COOK Report mail list, including 20 members who live in the Netherlands were queried as to why SURFnet had rarely been a subject for discussion. Was it irrelevant to everyday concerns such as the fiberization of Amsterdam? Or was it because it was just unknown? Or was it some unknown factor?

The replies were illuminating.

Rudolf van der Berg: see: http://internetthought.blogspot.com/
a Telecommunications consultant ICT consultant wrote:

"Hendrik Rood or Jaap van Till are probably much more versed in the ways of SURFnet than I am. I personally don't see SURFnet as such an enormous accomplishment. Given the size of the Netherlands and given the nature of the demands I wouldn't know how you would deal with interconnecting academic institutions differently. If anything what helped was that the Netherlands realized that academic networking was important earlier than others and by keeping that lead, make the Netherlands the default gateway to Europe. Stuff like lightpaths just follows from the technology and isn't really that flabbergasting. That said if you compare SURFnet with the rest of the world, we somehow seem to

have done right or better said much of the world seems to have done wrong. But let me explain why I feel it is such a natural thing to do and why I do believe SURFnet was involved in building these ideas.

"SURFnet meant there was always enough bandwidth..."
RUDOLF VAN DER BERG

I came online in 1994 at Twente University where my current group of friends organized as Datanet Drienerlo (DND) were rolling out a 10 Mbit/s Ethernet network on campus and 100 Mbit/s switched was available by 1998. At first this was an un-official non-university sanctioned network built with whatever was available and a good drill to get a cable from one dorm to the other. From 1994 Twente University paid and rolled out the network, but the students operated the network. The campus net was and still is hooked up to SURFnet with at first an ATM uplink, but later on GigE and 10 GigE. The students have been demanding customers ever since the Campusnet started. Much legal and illegal traffic was handled on the network, but for instance standard P2P failed to catch on as everybody shared a

couple of folders publicly within the campus and someone build a cool search interface to make it all accessible. *The network became an enormous breeding ground for new ideas and network use.*

Having joined that group in 1998/99 and becoming vice chairman of the campus network, I only learned about networking in the Ethernet/IP way of thinking. Traditional networking based on circuit-switched or smart network ideas have always felt erroneous to me. Whenever we played with new technologies (like streaming HD video on demand in 1999/2000) we did not need to look at the network and looked at the service instead. When a company asked us to test drive their HD video on demand server, some friends rebuild their two years of work within a day and then some tweaking for a week. There were lots of questions whether Linux or the machine could handle the task, but no one asked if the network could handle it. It was assumed it could despite all traffic of the students accessing movies across the network on hard disks on other PC's across campus which would compete with the streaming VOD for bandwidth.

DND is still active after more than 20 years and can now be credited to have been in-

volved with the establishment of the campus wireless network, my first employer the local internet exchange (NDIX), the Trent regional telecom network and I dare say even with some of the thinking behind FTTH in the Netherlands. For instance some of Reggefiber's people had an office in the Virtu Secure Webservices (now Equinix) building where the NDIX was located as well. We were discussing FTTH then and some of us are still involved with it now. SURFnet wasn't involved there, though connections to SURFnet were one of the reasons why customer owned regional fiber took off in the East of the Netherlands.

"...people just don't realize infrastructure is crucial until things go wrong..."
LEO PLUGGE

No matter which way you look at it. SURFnet wasn't a direct driving force behind much of what happened in the Netherlands, but it was involved or a known quantity in the background. SURFnet allowed a way of thinking that assumed there was always enough bandwidth available and it often helped out innovative parties with

that bandwidth. Once that way of thinking is with SURFnet as part of your telecom genetic make up, it becomes normal to think in GigE uplinks instead of T1's. I do believe that SURFnet allowed a whole generation of network nerds not just at Twente, but at every university who in their formative years in university knew just one thing: "The network can handle it".

Leo Plugge, Secretary of the SURF Advisory Committee: "It's interesting to see how SURF's infrastructure is taken for granted - or even as natural - by many users. As Rudolf wrote: "I wouldn't know how you would deal with interconnecting academic institutions differently", and students know "The network can handle it".

…government decision to treat it as an infrastructure of national importance…

Infrastructure is essential, but it tends to be hidden in daily life. Only when things go wrong people (users) become aware how crucial infrastructure is. As the University of Twente learned the

hard way after a fire in their computing center in November 2002. Or take New Orleans as an example. But since this was infrastructure, as Kees has pointed out "Together with BT, Cisco and the University SURFnet were able to install a complete new SURFnet node at another location on the campus, including a new GSR router and fiber connection, which became operational the next day after the fire." It's not unlike our water management infrastructure that is an example to the rest of the world, but is taken for granted by most Dutchmen. Without a tea party to underscore the importance of infrastructure, we run the risk of dozing off, as we did before the 1953 flooding, which was a very hard wake-up call indeed.

Is it a national asset? Yes. Our current IT-infrastructure may seem a natural thing to do, but that does not mean that it's not unique in its kind. Take our water defenses. The fact is, that the Dutch organizational and technical infrastructure arose naturally from villagers who organized themselves locally into Waterschappen since the 12th century. This was natural, but also unique in the world.

Having delegated the responsibility to solve the problem also means that those who

benefit become detached to the problem and the solution. As long as there is no major disaster/disruption/shortage, there is no need to interfere or worry. The same goes for electric power provisioning).

Luckily, we had a few visionary people come-up with the idea to create SURF in 1985 to realize SURFnet, and we now have the government decision to treat it as an infrastructure of national importance.

Herman Wagter (who is building Citynet in Amsterdam): As proud as we may be, there are more self-organized institutions in the world who have been long lasting, even longer than the Waterschappen.

The last Nobel Prize winner Elinor Ostrom has written about them in a commendable (very readable) book, Governing the Commons, that shows how classic economic science has ignored the potential for self-organization, assumed that the "tragedy of the commons" is a basic law (which it isn't).

"…it's no accident SURFnet matured into a major infrastructure."
VINT CERF

Jaap van Till: AMS-IX, SURF and SURFnet are in a sense cases of "Triumph of the Commons". Well understood of what resources/services are for shared interests (non combat) and definition of what not shared (competitive).   That NetherLight, GLIF, GLORIAD, and in a certain sense 'the national layer two SURFnet optical network' are not well known in the Netherlands is that these are big-pipes for a special group of big-science like nuclear data of CERN and long baseline radio astronomy.

These applications need so much bandwidth capacity (and have limited budgets as always) so they have been taken off the Internet connections. Otherwise the rest of the internet in Europe would have grinded to a halt during these trans-national scientific experiments. It was recognized long ago here that it is useless policy that by thinking (vertical) of new services you can stimulate network use. It does not work like that. Scalable network infrastructures (in horizontal layers) will boost invention of new applications by removing obstacles to growth.

Good roads create traffic. You cannot boost the growth of grass by pulling at its leaves. It is better to spread some infrastructural manure, water and sunlight.

Steve Wolff: However natural SURFnet may seem to its clients and customers, it has undeniably been from its very beginnings a technology leader in the worldwide NREN community. In addition - at least from an outsider's perspective - it has been a model of how to deal with its national government in providing services and avoiding the fate of those NRENs (National Research and Educational Networks) who are increasingly seen as competitors of unsubsidized private ISPs.

[**Editor's Note:** Steve Wolff explained to me later the same day in a voice conversation that far too many NRENs are seen by their university CIOs as means of getting subsidized broadband for web and email for students and faculty and have precious little to do with enabling applications not possible on commercial networks.]

One COOK Report member who responded was not a citizen of the Netherlands, but he does know something about SURFnet. He is the man many call "the father of The Internet" and currently serves as Chief Internet Evangelist for Google.

Vint Cerf: I have been following SURFnet and the SURF team since its origins and have worked closely, from time to time, with Kees Neggers and others there. Frankly, I have been extremely impressed by Kees and his colleagues. Their ability to be in the center of very high capacity networking activity and be reliable providers of super high speed platforms is a significant accomplishment and certainly put the Netherlands at the leading edge of academic networking. If their effort has now matured into a national infrastructure, it is no accident.

Exchanging a temporary resource for a sustainable knowledge infrastructure

A paper by the Netherlands Bureau for Economic Policy Analysis gives an excellent overview of Dutch budgeting philosophy and the thinking that undergirds the use of gas money for infrastructure. Readers will see that it ends with great uncertainty about the sustainability of the Dutch economic course. (The COOK Report does not share that uncertainty.)

From The Dutch Fiscal Framework: History, Current Practice and the Role of the CPB, Frits Bos, CPB Netherlands Bureau for Economic Policy Analysis --

http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1000121

From the Abstract: "According to the IMF and OECD, the Dutch fiscal framework is rather unique, and its design and implementation are highly recommendable.. This paper discusses this framework. ... Since 1814, the official notion of a balanced budget has changed substantially over time. First, when debt was excessive, it imposed the redemption of loans. Later, a golden rule of finance was introduced, allowing new loans for 'productive' expenditure.

At the end of the nineteenth century, the prominent Dutch economist and politician Pierson stressed that each generation should bear its own burden and should not leave excessive debt for the next generations. After the Second World War, the classic view on the government was replaced by a macro- economic view: the budget of the state was presented as part of a set of national accounts on the Dutch economy. Since then, the CPB, being an independent institute, provides the official estimates on the macro-economic developments. Directly after the Second World War, this new macro-view was combined with a strict budgetary control ...

Since about 2000, a forward looking view on Dutch public finance has become dominant: Dutch public finance should be sustainable in view of the net extra costs of aging and the falling revenues from natural gas." **pp. 9–10**

The Economic Structure Improvement Fund (FES) was established in 1993. Government investments in infrastructure had fallen from about 3% GDP in 1970 to 1.5% GDP in 1993. By earmarking via the FES about 40% of the natural gas revenues for financing "additional investments of national significance", the structure of the Dutch economy should be improved. ... Since 1993, the FES has disbursed more than 31 billion euro. In the beginning, the FES-investments mainly focused on transport and mobility, e.g. roads, railway-tracks and channels. However, now also expenditure on knowledge, innovation and the environment are financed via the FES. **p. 33**

> *"Dutch gas reserves will be gone in 25 years. What then?"*
> *FRITS BOS*

The major assets of the Dutch government are the natural gas stock, the fixed capital stock and the financial assets. The discounted value of the natural gas stock was 90% GDP in 1970. At present, it has declined to 20% of GDP. **p. 34**

Mid 1990's, Dutch politicians explicitly addressed the issue of sustainability by creating two funds: the FES-fund and the old age state pensions-fund. These should help to ensure sustainability of Dutch public finance in view of the exhaustion of natural resources and the expected rise in old age state pensions due to aging. However, the solutions offered were only formal solutions: 40% of the natural gas revenues was to be used for financing FES-investments. The motto was to turn underground assets into assets above ground. In particular when cost-benefit analysis for FES-investments was not obligatory (see text-box on FES), there was no guarantee that this results in a higher return than alternative options, e.g. extra expenditure on education or extra reduction of public debt. As a consequence, the FES-fund is important for changing the composition of public expenditure, but its contribution to sustainability is not clear. **p. 35**

Under unchanged polices, the aging population will lead to a sharp and structural increase in public expenditure, in par-

ticular on state pensions and health care. Government revenue from taxes on funded pensions will also increase, but not enough to cover the extra expenditure and the falling revenues from natural gas. As a consequence, in the long run without policy adjustments public debt will explode and Dutch public finance will be out of control." **p. 36**

We end our tour of the Dutch Fiscal Framework paper with this quote: "In about 25 years, Dutch natural gas reserves are expected to be exhausted." **p. 37**

The most surprising aspect of the above paper is how oblivious its author appears to be regarding the critical role digital technology, and especially networked digital technology, plays in redefining how the world will work and who will succeed from this time forward. The late Hans Rosenberg understood this, as does everyone at SURF and at other institutions in the Netherlands ICT ecosystem. It's why the government in the Netherlands has chosen to make this investment.

The genius of the Dutch national character is, as we wrote earlier, the often-demonstrated ability to work together to find and execute strategies that win significant competitive advantage for

the Netherlands. Establishing a benchmark ICT technology infrastructure is in that tradition. It has given the Netherlands "the high ground" in this new global environment where the world of atoms and the world of bits are rapidly becoming one. (Giving the Netherlands "the high ground" is a nice achievement, metaphorically speaking, given high ground is something the country lacks.)

The Netherlands is advancing towards gaining "first mover advantage" in the world of research and collaboration networks.
http://en.wikipedia.org/wiki/First-mover_advantage .

…gaining first mover advantage gives the Netherlands the high ground.

Creating a world-leading knowledge infrastructure, that is further empowered by its open philosophy, is likely to make it more nimble and successful than the tendency of other EU initiatives, which continue to favor large scale top down projects. It seems extremely likely that the SURFnet7, GigaPort3 and e-Science Center initiatives will prosper. They will continue to create a foundation for an

environment designed to enable scientific and technology success. This is the kind of economic growth that, with pragmatic fiscal policies and competent governance, could eventually replace the income generated by millions of years of geologic history.

The July 2007 study concludes with the warning that given that gas reserves were finite (25 year estimated life time), it was important to transfer infrastructure support from FES gas fund basis to a structural basis meaning part of the normal government budget funded by taxpayers and the NWO – that is the Dutch Science Foundation. The section below, called the Dutch Road Map, provides to find out what happened.

The Evolution of Knowledge Infrastructure as a Generally Accepted Goal

One of the key questions is how does a physical infrastructure become a knowledge infrastructure? Kees Neggers says the process began in 1994 with documentation that formalized the goal. This initiative was refreshed and strengthened in 1997 and 1998. The Dutch were "leading the duck", once again taking the First Mover Position, as the EU didn't be-

gin talking about having a knowledge infrastructure until 2000. The EU initiative resulted in the Lisbon Protocols.

According to **Wikipedia**, the Lisbon Strategy, also known as the Lisbon Agenda or Lisbon Process, is an action and development plan for the European Union. Its aim is to make the EU "the most dynamic and competitive knowledge-based economy in the world capable of sustainable economic growth with more and better jobs and greater social cohesion, and respect for the environment by 2010".[1] It was set out by the European Council in Lisbon in March 2000. Contemporary key thinkers on whose works the Lisbon Strategy is based and/or who were involved in its creation include Maria João Rodrigues, Christopher Freeman, Bengt-Åke Lundvall, Luc Soete, Carlota Perez, Manuel Castells, Giovanni Dosi, and Richard Nelson. Key concepts of the Lisbon Strategy include those of the knowledge economy, innovation, techno-economic paradigms, technology governance, and the "open method of coordination" (OMC). http://en.wikipedia.org/wiki/Lisbon_Strategy. But the Lisbon Strategy was viewed as a failure long before 2010.

Meanwhile, the Netherlands focused on creating a digital commons in the form of ac-

tual networks, services and tools that could be shared and could attract other projects and communities within broadly accepted national goals.

SenterNovem Turning Policy into Reality
http://www.senternovem.nl/english/

The COOK Report is especially interested in SenterNovem, an agency of the Dutch Ministry of Economic Affairs. It's website states: "We promote sustainable development and innovation, both within the Netherlands and abroad. We aim to achieve tangible results that have a positive effect on the economy and on society as a whole. Our core competence is converting government policy into reality. On behalf of the Dutch government we implement policy regarding:
--Innovation
--Energy and Climate Change
--Environment and Spatial Planning "

…if the Netherlands wants to be truly competitive, we need to be innovative we need to excel…
WWW.SENTERNOVEM.NL

It isn't fancy, but visitors can quickly click through and find

a number of real, active programs to facilitate the use of the evolving Knowledge Infrastructure in the Netherlands.

COOK Report member **Rudolf van der Berg** explained what SenterNovem did and how it actually worked. He wrote on the list: The short answer is that the Ministries make the policy and SenterNovem is just the factory that processes all the subsidy and permit applications that are the result of this policy.

In the Netherlands there is a (strict) separation between formulating policy and executing policy. The Ministry of Economic Affairs (COOK Report list member **Joost van der Vleuten**) is currently formulating a policy on how to get broadband to all. Let's say for arguments sake that Joost actually got €7 billion to roll out FTTH to every home.

[Editor's Note: Of course not likely given the current economic situation.]

After this policy is approved by the Minister or if it goes beyond her powers, by parliament, SenterNovem will most likely be contracted to execute the policy (distribute the money in this case)."

Van der Berg continues: "SenterNovem will look at the policy as already defined and

**Motion Submitted by
Marjet van Zuijlen
6 November 1997.**

Motion The Parliament, heard the discussion,

whereas, The Netherlands has the ambition to be in the front line of the electronic highway;

Considering that these developments until now have seen a slow start;

whereas, the government is generously investing in physical infrastructure;

asks the Government to encourage that for the destination of the so called ICES-funds a balanced approach is chosen in which there is also scope for strengthening the knowledge infrastructure, such as investments in SURFnet4 and Amsterdam as an international data communications hub,

and proceeds to the order of the day

OECD are done by the Ministry. OPTA quite often has an input in these areas, though not always. For instance Lawful Interception or ICANN and the OECD are generally not of OPTA's concern at all.

COOK Report: An extremely useful SenterNovem page that explains the development of investment in Knowledge infrastructure is found here http://www.senternovem.nl/bsik/algemeen/achtergrondinfo/index.asp
The website is in Dutch, but Google translation services provides a rough English translation.

then will work with the Ministry in further defining the nitty gritty of the policy. That is if the policy says that the money should go to the consumer and the consumer can then pick a network of choice, SenterNovem will have to decide whether it is permissable for FTTH-networks to combine the applications of thousands of individuals in a region to get the subsidy directly. SenterNovem will also decide how they want to verify whether the money has been used according to the rules etc."

"There is of course not a strict line between policy and implementation... it's a grey area at best. Depending on the area in which you are working SenterNovem will be more or less involved in the defining of the actual policy. For instance when it comes to renewable energy it has much influence as it employs

specialists in the area whereas the ministry consists more of generalists. When it comes to supporting the Creative Sector in The Netherlands it acts more as a facilitator, bringing the right people together. With certain simpler subsidies/permits for SMEs, it may be just a subsidy/permit factory, where the employees add little to the actual policy making."

The same division can be seen between the Ministry of Economic Affairs and the Dutch Telecom Regulator OPTA. Unlike the FCC or Ofcom, OPTA doesn't write it's own policy. OPTA has to execute what the minister has decided. In practice of course there are many shades of grey there. The Dutch input to the EU's Telecoms Framework, E.164 and E.212 numbering policy, policies on lawful interception or the interactions with ICANN and the

That webside focuses on the importance of Global Knowledge, identifying it as the key competitive factor. "The Dutch economy is increasingly knowledge intensive. Nevertheless, our economy in terms of knowledge is not up to the demands of our time. If the Netherlands wants to be truly competitive, we must be innovative, and dare to excel. This calls for drastic measures. It means that more research should now focus on innovation and/or social need. It also forces more effective deployment of resources and research, for example by joining forces of the involved parties. The Dutch government is aware and has taken steps to exploit innovation opportunities in the Netherlands in order to

propel it to the top in Europe."

Investing Together

Senternovem encourages co-investment and public/private partnerships, similar to the innovative infrastructure-sharing partnership recently initiated by SURFnet with Philips. (The Phillips partnership is discussed later on in Chapter VIII How A Progressive ICT Infrastructure Benefits The Economy".)

The SenterNovem website explains, "The government can not solve everything on its own. It is important that companies, research institutions and governments can invest in the development, dissemination and utilization of advanced knowledge. The Interdepartmental Committee on Enhancing Economic Structure (ICES) was established to promote investment projects that enhance the economic structure of the Netherlands. . . . . Because of the importance of investing in knowledge (KIS) the working group ICES/KIS was created. This group focuses specifically on projects in knowledge infrastructure. Under the leadership of the group, two knowledge infrastructure programs were established in 1994 and 1998 under the names ICES/KIS-1 and ICES/KIS-2. There was respectively 113 and 211 million Euros

available for these programs."

The vote to fund a knowledge infrastructure was virtually unanimous...

SenterNovem's website provides an historical timeline. The focus on economic infrastructure began with creation in 1993 of the FES fund or Fonds Economische Structuurversterking, (economic structure enhancing fund.) And in 1994 with ICES/KIS-1. But the parliament was of the opinion that not enough knowledge infrastructure projects were granted following the 1994 initiation of ICES/KIS1.

Then in November, 1997, something happened that opened the door to the future for creating a true national

knowledge infrastructure for the Netherlands. A 30-year-old history and political science major and Labor Party member named Marjet van Zuijlen went to parliament. http://nl.wikipedia.org/wiki/Marjet_van_Zuijlen Previously Ms. Van Zuijlen had worked several years as a recruiter in the personnel department of KPN. But on this day she was focused on the public interest. Her motion asked for the government to commit to faster implementation of the change towards knowledge infrastructure support, to finance entities like SURFnet4 to strengthen the national knowledge infrastructure and to position "Amsterdam as an international data communications hub…" The motion carried almost unanimously. It was the turning point that meant projects like GigaPort and VLAM (Virtual e-Science Laboratory Amsterdam) got funded out of FES at the end of the nineties.

Since 1993 Five ICES/KIS Calls Have Been Issued

1994 ICES/KIS-1 for 113 M Euro

1998 ICES/KIS-2 for 211 M Euro

2003 ICES/KIS-3 (called Bsik) for 800 M Euro

2008 FES round 279 M Euro

2009 FES round 500 M Euro

SenterNovem says this "speed up the 1998 call which was really focused on knowledge infrastructure and resulted in the release of 211 million euro for ICES/KIS-2." [Note the first GigaPort project in 1999 was funded out of ICES/KIS-2.]

As ICES/KIS-2 finished out in 2002, the umbrella for a follow-on program was called Bsik or Decision on Subsidies for the Investment in knowledge.   On November 28, 2003 the government made an amount of 800 million Euros available for knowledge based projects as part of the so-called "Third boost investment in knowledge" (ICES/KIS-3). The 800 million comes from the Economic Structure Enhancing Fund.

"There are seven ministries involved in the program, namely of Education, Culture and Science, Agriculture, Nature and Food Quality, Transport, Housing, Spatial Planning and Environment, Finance, Economic Affairs and Health, Welfare and Sport."

Bsik is to promote cooperation between research and enterprise development and thus leading to quality networks in the knowledge infrastructure. Within these networks, research is conducted that addresses the societal need for knowledge.

"Netherlands has opted for a limited number of strategic knowledge and the pooling of financial resources and expertise. All under the motto "innovation, focus and strength. The strategic knowledge are chosen based on ideas and needs from the field. . . . . This resulted in five themes: High quality space, ICT, Sustainable Innovation System, Micro and nanotechnology and Health, nutrition, gene and biotechnology breakthroughs."

Bob Hertzberger and Kees Neggers constructed a table for this report that illustrates the key events and players in the FES fund infrastructure policy process. It can be seen on the next page.

- The process starts in the left upper corner with the preparation of a new call. ICES/KIS2 for example is such a call. The government decides the amount of money available and the rules for judging the call. These rules and processes do not stay uniform from call to call.

- After Parliament approval, a call is published by SenterNovem. Public private partnerships will then be formed to react to the call and proposals written and vetted within each requisite industry and agency and its academic counterparts.

- Proposals will then be sent to the Committee of Wise men and to the Netherlands Bureau for Economic Policy Analysis each of which scores them and recommends some for funding while turning down others.

Government finally decides on the winners. The appropriate ministry in conjunction with SenterNovem will then oversee the execution of the projects.

- As a part of the funding rounds Government may decide that only one proposal from a strategic area is allowed, forcing the area to really join forces in a collaborative way. In the 2009 FES round, for example, ICTRegie was asked to prepare a single proposal for the ICT area.

The Dutch ICT Roadmap to a Generic Knowledge Infrastructure

The proposal progress mid 2008 through 2009:

Since the mid 1990s both the Netherlands and the European Union have set up planning commissions for infrastructure — and especially for research infrastructure. The Dutch Cabinet has the final vote on infrastructure projects. But the Cabinet must decide according to a struc-

FES Knowledge Infrastructure Proposal Procedure
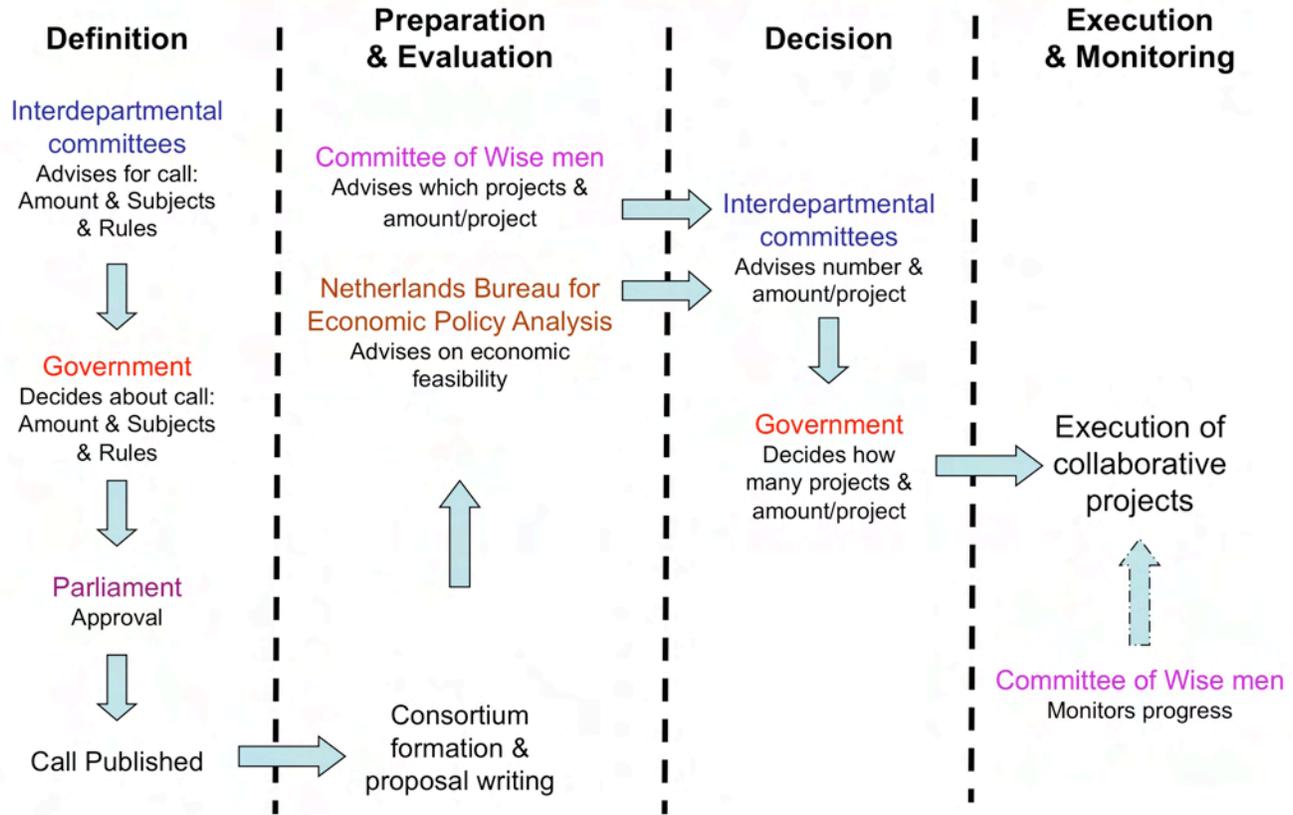
COOK Report: The procedure for deciding how to use infrastructure funds fascinated me and as I tried to understand how it worked the vision of a kind of flow chart occurred. so thanks to Bob Hertzberger and Kees Neggers for designing the above chart. Follow the arrows. I wish hat the US government would do this.

tured and well thought-out process.

FES rounds are made on a regular basis for funding with specified blocks of gas money. They are reviewed by the Commissie van Wijzen (the Committee of Wisemen) for technical soundness and by the Netherlands Bureau for Economic Policy Analysis

(CPB) which renders judgment on the estimated economic benefits of the proposals. The two findings and recommended actions are then sent to the Cabinet for the final funding decision. (See the diagram above.)

The July 2007 CPB Report cited above caused the Minister of Education, Culture and

Science (OCW) to set up the Dutch Committee on a National Roadmap for Large-scale Research Facilities. This became known as the Van Velzen Commission, named for its Chair Wim van Velzen. Mr van Velzen is a very senior European policy advisor in focusing on technology assessment and science. He has also very recently been

named Chair of the Committee of Wisemen. He is a member of the Christian Democratic Party (CDA). Previousy he served as Chairman of CDA (1987-1994), as Senator in the Dutch Parliament (1987-1995), and as a member in the EU parliament (1994-2004).

Wim van Velzen

Also in 2007 ICTRegie began responding to a request from the ministries of OCW and Economic Affairs for an evaluation of the entire ICT Infrastructure for research in the Netherlands. In an interim report in December 2007, the Van Velzen Commission had indicated that high quality infrastructure for research was an essential prerequisite for hosting any of the large-scale ESFRI facilities in The Netherlands, or to allow Dutch researchers to use these facilities elsewhere. ICTRegie stressed the importance of an excellent ICT Infrastructure in a February 2008 report as well.

Instruments, computing facilities and storage are becoming generic resources.
SURF ICT Roadmap

[**Editor's Note:** The European Strategy Forum on Research Infrastructures (ES-FRI) established in April 2002. The role of the European Strategy Forum on Research Infrastructures (ES-FRI) is to support a coherent approach to policy-making on research infrastructures in Europe, and to act as an incubator for international negotiations about concrete initiatives. In particular, ESFRI is preparing a European Roadmap for new research infrastructures of pan-European interest.
http://www.lifewatch.eu/index.php?option=com_content&view=article&id=5&Itemid=5

There is increasing broad-based interest in many EU countries today in increasing the visibility and viability of European ideas in technology and culture.]

In responding to Van Velzen and ICTRegie, OCW in effect said: "We want you to look at the long term sustainability of this infrastructure. We want you to figure out, given the expected gas fund deple-

tion over the next 25 years, how to make the knowledge fund infrastructure structurally sustainable — i.e. to make it supportable by the taxpayer and from regular government budgets.   The Van Velzen committee and ICTRegie invited SURFnet, the VL-e project and other stakeholders to explain what they had done in the past, hoped to do in the future, and to estimate how much it would cost.

Consequently SURF responded on July 8, 2008, with a ICT Roadmap proposal that described what a National ICT Research Infrastructure for the Netherlands would look like. It was prepared collaboratively with inputs from SURFnet, Netherlands Computing facility (NCF), BigGrid (see later in VL-e interview) and VL-e with the help of others such as SARA and Netherlands BioInformatics Center (NBIC). Here are some of the data points from the SURF response.

"Modern research is increasingly a global effort, based on international cooperation and resource sharing. Instruments, computing facilities and storage are becoming generic resources, which can be combined through advanced networks to provide services to an increasingly diverse research community.

"The exchange, combination and integration of data and information has become possible as computer and storage facilities have been connected via very fast and high bandwidth networks. The resulting distributed computer systems are developing into a single ICT research infrastructure, consisting of hardware (network, computers and data storage facilities) and software (web and grid middleware) harnessing the resources. By enriching this infrastructure with adequate tools, which provide services for e-science, a generic instrumentation for modern research emerges which is far more effective than anything the individual research disciplines could develop themselves.

"This proposal concentrates on creating and maintaining an advanced ICT research infrastructure in the Netherlands. It includes networks, computing, vizualisation and storage hardware, as well as the middleware and generic services needed to enable modern research. Given the importance of international cooperation in modern science, the infrastructure will be connected to other initiatives worldwide.

"The project will build upon existing facilities and knowledge where possible. In particular, it will build on the results from the projects VL-e, BiG Grid, GigaPort and GigaPort Next Generation, NBIC and MultimediaN, as well as from the NWO-funded research programs for computational science and computational life sciences.

"To implement the infrastructure we intend to structure all necessary activities into a coherent program. Such a program will include the following activities: Acquiring, maintaining and operating equipment and software licenses; Software hardening and support; R&D to enable efficient usage of the infrastructure and to facilitate new communities."

The proposal to the Roadmap Committee had to discuss the criteria shown in the following text box below. In doing so they created the inventory summarized in the in Chapter VI Growing e-science Domains.

Funding Criteria

- **Science case**: potential for scientific breakthroughs;
- **talent case**: potential for brain gain, Netherlands has to remain an attractive and challenging place to work;
- **partnership**: opportunities for collaboration and synergy;
- i**nnovation case**: interest for society and industry, magnet for new knowledge;
- **business case**: financial soundness;
- t**echnical case**: technical feasibility/challenge.

And
- **focus for the Netherlands**, either in existing expertise or desired expertise;
- **critical mass**, availability of sufficient Dutch qualified top-researchers interested in this facility;
- **embedding**; the facility should be embedded in Dutch and international larger research structures or stimulate further collaborations and concentration of efforts;
- **proven will** to collaborate, including commitment for own financial contributions;
- **connection to societal developments** and challenges, like water, energy, health care, security etc.

III. The Direction of ICT Infrastructure in the Netherlands in Early 2010

An interview with Wim Liebrand, Kees Neggers and Cees de Laat

The ICT Infrastructure for the Netherlands continues to evolve, mature and improve it's ecosystem. And unlike that classic productivity graph for Moore's Law – with it's single curve arching upwards – a collaborative, network-based, re-usable infrastructure evolves, matures and improves in many directions at once.

On Tuesday, the 17th of November, 2009, many of the leading figures in the ICT Infrastructure for research in the Netherlands sat down with The COOK Report at the Supercomputing 09 meeting in Portland, Oregon, to provide a progress report on the state of Netherland's infrastructure as a new decade begins.

The participants included:

Wim Liebrand, Director of SURF

Wim Liebrand has been the Director of the SURF foundation since September 2001. He explained that SURFnet is an important division of SURF for which SURF is responsible. There is also a division that deals with software licensing for staff and students in research and higher education institutes and next year SURF will start a division that deals with shared services for students' admissions and enrollments in Higher Education. After the implementation of the ICTRegie advice we will also be responsible for the computing and storage facilities and will be supporting a new e-Science Research Institute.

Hans Dijkman is responsible for E-science at the University of Amsterdam. He has a PHD in Chemistry and an MBA. He has 20 years experience as ICT manager of ICT departments from research and education institutions. His Management philosophy: Mathieu Weggeman, Òn management of professionals? Don't!

Bob Hertzberger is director of the VL-e project as well as one of the directors of Big Grid and till recently adjunct

director of Netherlands Bio-Informatics Center (NBIC). As VL-e director he was responsible for developing the functional model that together with the methodology of multidisciplinary collaboration form the basis of the project.

Kees Neggers is one of the founders of SURFnet in the Netherlands and has been a Managing Director there since 1988. Since the mid-eighties Neggers has been on the forefront of both the technical and organizational evolution of the internet and research networking developments worldwide and involved as initiator and Board member in many international networking related organizations and initiatives like RARE,

Hans Dijkman, Director e-Science University of Amsterdam

TERENA, Ebone, Internet So-
ciety, RIPE NCC and most
recently GLIF, the Global
Lambda Integrated Facility.

Cees de Laat is professor
and chair of the System and
Network Engineering group at
the University of Amsterdam.
Research in his group in-
cludes optical /switched
Internet for data-intensive
TeraScale e-science applica-
tions, Authorization, Security
and Privacy in distributed
systems and Semantic web to
describe e-Infrastructure. He
serves as board member of
ISOC.nl, chairs GridForum.nl,
is co-founder of CineGrid.org
and the Global Lambda Inte-
grated Facility (GLIF).
http://ext.delaat.net/

The Technology Research Ecosystem of The Netherlands

COOK Report: No other
country has put together
such an extraordinary eco-
system of optical infrastruc-
ture and leading-edge re-
search and industrial tech-
nology transfer applications
as has the Netherlands. *This
is a true 21st century infra-
structure* that is, in a way,
equivalent to the remarkable
17th century infrastructure of
windmills and canals. That
was a cooperative pragmatic
environment that enabled the
invention of the mechanical
sawmill in the 17th century.
In turn that led to the defeat

of the Spanish Armada by
making it possible to build
more and cheaper ships.

you have built the 21st century infrastructure

It is impressive that the
Netherlands managed to
gather together the neces-
sary academic, political, sci-
entific and research funding
entities in the Netherlands —
to produce what appears to
be a unique approach to the
careful coordination and
planning with maximum cost-
effectiveness of public and
private funds — resulting in
leading-edge research and
technology transfer. In many
respects, you now have the
most advanced optical net-
work in the world.

So where do you go from
here? How did you get to this
point? What have you ac-
complished and where are
you are headed in the near-
to-midterm future?

Neggers: In December,
2008, the ICTRegie published
its recommendations for the
future of the national ICT in-
frastructure. The government
took that advice into consid-
eration, and by May, 2009,
they had decided the ICTRe-
gie advice was good and
wanted it to be followed up.
Consequently they asked
SURF to make an implemen-

tation plan for it.

Liebrand: When you talked
to Kees Neggers a year ago,
the ICTRegie Report was still
under construction. It took a
couple of months of review
before it was fully supported
by the government — that is,
the Ministry of Education and
the Ministry of Economic Af-
fairs. So they together - re-
search and industry and edu-
cation –supported the ICTRe-
gie document. They asked
SURF to work with these enti-
ties to build a transparent
plan for an eco-system for
technology transfer from ba-
sic research.

The total financial picture has
not yet been realized as the
result of the impact of the
global economic crisis, but
the good thing is that every
one is collaborating and we
are finding piece-by-piece
financial support with which
to implement the plan by
means of cooperation be-
tween industry government
and the research institutes.

[**Editor's Note**: The recom-
mended yearly budget for the
implementation of the IC-
TRegie advice is 62.5 million
Euro, split over the various
infrastructure items as
documented in chapter VIII,
page 26 of the ICTRegie ad-
visory from December 2008.
The Networking component
of this is 10 million Euro per
year. And it is already se-

cured for 5 years for Giga-Port3.]

COOK Report: Is it the shared vision of all participants that has provided the critical difference in the Netherlands?

…the network is a condition sine qua non…

Neggers: Yes I think you can say there is among the stakeholders a shared vision that was achieved by extensive consultation among them before the report was published. Consequently we have a consensus among the stakeholders that this was indeed good advice. Now at that point the government needed to take it on board and make a financial commitment to implement it.

As Wim has pointed out the government said: SURF – you will be the umbrella organization to govern this whole eco-system. Now give us a detailed plan as to how you will do this. And until we have seen your plan and decide that it is a good and workable plan, we won't commit the money. This is the position we are in at the moment. The financing of the network innovation however is already committed. Everyone was convinced that a

network is a condition sine qua non in all scenario's.

Collaboration and the E-science Paradigm

COOK Report: Are you able to talk about the plans?

Liebrand: The plans themselves are described in the advisory report. And together with the Dutch NSF (which we call NWO) we are working to build an e-science research platform… which is a combination of fundamental generic research, combined with the infrastructure. That in turn is combined with the network.

SURF You will be the umbrella organization to govern this whole eco-system…

COOK Report: Is this what Bob Hertzberger is doing?

De Laat: Yes, Bob Hertzberger is talking about the software on top of the infrastructure. The understanding in the Netherlands is that software is actually infrastructure. Without the middleware to enable the infrastructure for applications the infrastructure itself is

useless.

COOK Report: So when you say software, you are talking about "middleware"?

De Laat: Yes, the generic part and the application domain specific part - both are middleware.

Liebrand: E-science has always been difficult to label, but we understand it in the broad definition provided by Bob Hertzberger later in your report. However, e-science is also related to computational science, to simulation, and to the grid computing and visualization infrastructure. Take all of these together and you will get what we consider to be the e-science paradigm.

De Laat: It is a new way of doing science.

COOK Report: It is the dikes and canals of the twenty first century.

Dijkman: Not just the dikes and canals. It is also the ships.

De Laat: Of course, you also need the correct people to steer those ships. It needs knowledge and tools by which the infrastructure can be used.

Software is way more expensive than hardware. For hardware you can get quite a discount for an expensive

cluster or supercomputer. But when, for example, you consider the man-hours that have gone into creating Globus, (a particular implementation of the Grid middleware layer) you understand that a floppy with a Globus installation is very expensive. The real cost for creating the correct software is typically underestimated.

Globus is open source. Its libraries can be freely downloaded. But the manpower required to manage and maintain these generic environments of the middleware, which make the grid and clouds possible, is always under-estimated. And the middleware *must* be kept up to date so that the biologist, for example, can use the grid without having to become a computer scientist.

The real cost for creating the correct software is typically underestimated.

Neggers: This is very similar to the physical network itself. No single institute could build the network to provide necessary connectivity. This is why *you need cooperation and investments up front to create the infrastructure.* We have done this for the network. And then

we have to be careful that what we build works… and not just inside the Netherlands. It has to work throughout the global research community. That is one reason we have all come here to Portland, to make sure what we are building is available globally. The same is true for this grid and e-science middleware layer. No single institute – not even a smart guy like Cees de Laat can do it on his own.

This is why the government of the Netherlands is interested. Because they need to provide the lubrication that keeps things going forward.

COOK Report: What is the role of GLIF in all this?

Neggers: We started GLIF because… we were aware from the start… that if we built something that works between Utrecht and Amsterdam, but is not connected to the rest of the world in a viable fashion, it is really useless. This is why we peered with Tom DeFanti and Maxine Brown at STARTAP, and then StarLight in Chicago, from day one in 2002. At that moment it achieved a critical mass globally.

De Laat: Of course that first lambda connection between Amsterdam and Chicago didn't work just like that. We had to research why it did not perform and when we

found out, we published our results in scientific and engineering journals. It was, and often still is, not trivial.

COOK Report: Was the amount of data so huge that your buffers filled up while waiting for an Ack?

De Laat: Yes – we had to re-engineer the transport protocol stack. We get repetitions of all these kinds of problems as we increase our bandwidth (1 - 10 - 40 - 100 Gbit/s). I can point out some rather common problems in networks which pop up again and again every two or three years. For example, in some cases UDP traffic stops working after some time on layer 2 engineered lightpaths. Why? Because UDP is unidirectional, and your Ethernet switches forget where the destination is. And they start flooding.

We had to reverse engineer the hardware in the link to Chicago. With lambdas and lightpaths, you really open up the infrastructure to the applications, which become very sensitive to the properties of the hardware if you remove the router layer. We could really measure those properties, the buffers and bottlenecks and so on… and given the round trip time, we could predict how certain links would perform with certain protocol stacks and applications. I had to go back

into Cisco… and up to Chambers-minus-one… to get their attention to the limitations in the hardware.

Neggers: The initial response from Cisco was interesting. They said, you are using our equipment incorrectly. This is a metro box, and you are using it wide area.

De Laat: Then we proved

At the Netherlands booth SURFnet was demonstrating how easy it is to upgrade from 10 to 100Gbit/s second in its existing Nortel Common Photonic Layer equipment. Two SURFnet6 nodes were installed interconnected by an artificially long fiber of 1120 km. See the picture on p. 37 below for a close up of the fiber spools seven of which were linked together.

that even when the box is used in the metro, but services some connections that go on long distances, the problematic effects would still be there.

COOK Report: In 2002 you were addressing the recurring problems of the network and the middleware. Did you find a similar set that affected the high performance computers?

Liebrand: Yes. But viewed from a slightly different perspective. There are so many cores – as in multi-core processor -- that we have to think about how to use them properly. The old ways won't work anymore.

Return on Investment in the High Performance Infrastructure

De Laat: There are indeed similar issues in high performance computers. There is an accretion of knowledge that is built up, and it needs to be maintained as the infrastructure grows. One of the tasks for SURFnet and for e-science is that the knowledge base gained from using this extreme infrastructure must be maintained.

COOK Report: What do you mean when you talk about SURF providing the "oil" that lubricates this ecosystem?

Liebrand: Actually it is a rather unique phenomenon because the cooperation within our "umbrella environment" has existed now for some 25 years. We began in 1986 with SURFnet version 1.0 at 9.6 kilobits per-second.

The collaboration model is to build a pretty advanced infrastructure, with innovation money from government and industry, and sometimes from participating partners. And then you give that to the users as a production infrastructure. Meanwhile, as the users are gaining experience, you start with a *second* round of innovation… that builds an even more advanced infrastructure. We innovate successive networks. So with SURFnet6 mature, we are getting ready to launch SURFnet7 soon.

Neggers: It is interesting to note that the Government innovation grant which enabled us to build the first SURFnet network was roughly the same amount as the grant that enabled SURFnet5 with 10 gig lambdas and SURFnet6 as a full hybrid optical and packet switching network.

Liebrand: More bang for
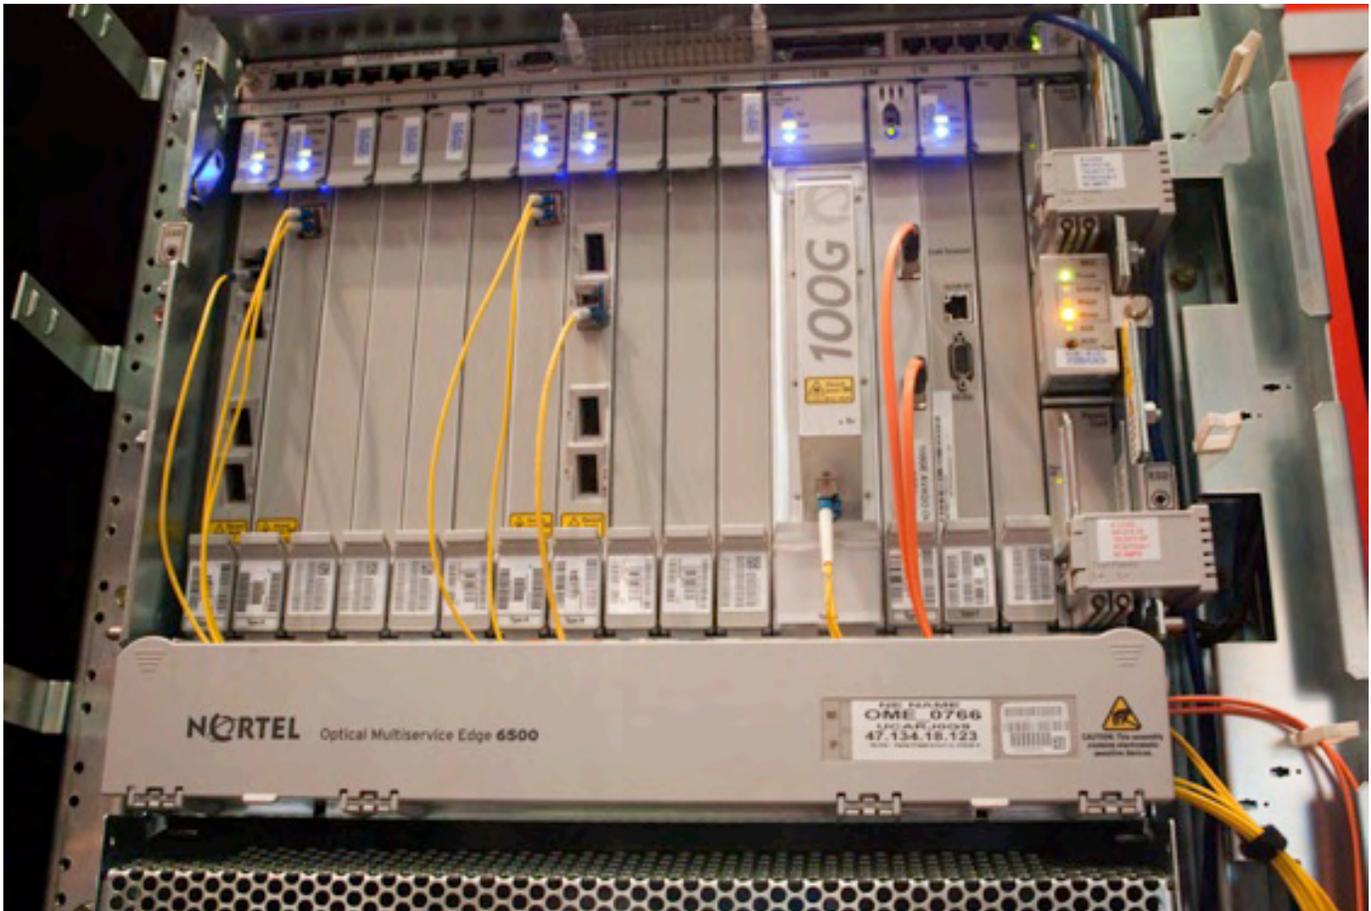
This picture shows a close up of one of the OME6500 systems in the racks at the booth. The only thing needed is two free slots in an OME6500 to put in a 100Gbit/s transponder. The 100G links operate on the existing 10 Gbit/s optical wavelengths. In July 2009 Nortel and SURFnet successfully demonstrated that this works in its live network over a distance of 1244 kilometers on its link between Amsterdam and Hamburg.
More information can be found on:
http://sc09.delaat.net/
http://www2.nortel.com/go/news_detail.jsp?cat_id=-8055&oid=100257448&locale=en-US&lcid=-1
If you are interested in a YouTube movie of the demo go to:
http://www.youtube.com/watch?v=P5jVGpjbdkQ
For a more technical demo in the Nortel labs go to:
http://www.youtube.com/watch?v=zEUDiRWNmII&feature=related
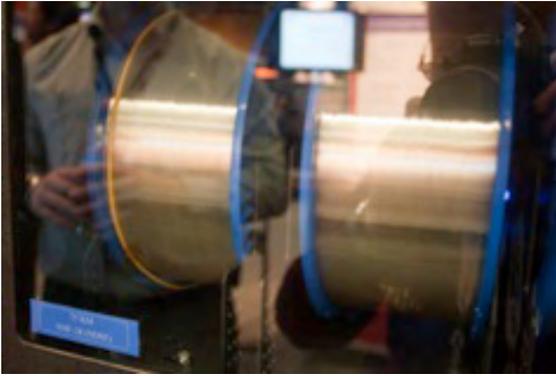
Close-up of the fiber spools. There were 7 spools of 80 km in one direction for a total of 560km with a loopback of the 100G wavelength for an additional 560 km of fiber for a total of 1120km between the two optical nodes for the 100G demo on the Netherlands booth at SC09 in Portland. The nodes were less than 2 meters apart on the show floor shown in the picture on page 35 above

your buck. This is Moore's Law applied to innovation networks.

De Laat: Just for fun I have calculated very roughly the increase in capacity for the Netherlands from SURFnet4 through SURFnet6. I multiplied the number of kilometers, first of copper, then of glass, times the capacity. The growth over 14 years is about 10 orders of magnitude. Actually, these increases are much faster than Moore's law!

Liebrand: It is difficult to do a precise calculation, but for each Euro invested, the payback to the economy is between four and five Euros. This is my gut level guesstimate.   Maybe Kees knows more advanced models?

Neggers: For SURFNet5, 6, and 7 our plans were evaluated by economists from the

Netherlands Bureau for Economic policy Analysis (CPB). The government requested this as part of the policy procedure for proposals funded under our knowledge infrastructure funds *(as described on pages 14 to 19 above).* They did econometric studies beginning at the end of the 90s. In the 90s investment in these technologies was regarded as a no-brainer. But after the dot com crash, it was no longer so easy to defend investments.

We had a difficult time convincing them to finance our innovation again out of these infrastructure funds. Initially they saw no need for a new investment in research networking or for tapping these special funds again. By that point they considered the network to be nothing more than a normal tool for the researchers. Consequently, it should be paid for only by the government entities that supported research.   Their initial conclusion was that these entities should pay for the network from their regular budgets. They no longer saw investing in the network as infrastructure as having a broader national impact.

We then asked for and re-

ceived supporting letters from stakeholders outside the direct Research and Education Community — who already had seen the positive effects of our activities for society as a whole. This, together with a further explanation from our team, convinced the CPB economists that what we did was important. And they endorsed our plans again.   Also the so-called Commissie van Wijzen, "The Wisemen", which at that time was the final advisor to the Government, also recognized once more the importance of our efforts for the entire economy.

NetherLight and AMS-IX are both real examples of unanticipated payoffs from the SURFnet projects.

In 2008 things became a bit easier again. Just before the government evaluation of our SURFnet7 and GigaPort3 plans, an evaluation by an independent party was published.   It focused on the whole ICES/KIS-2 program that financed the first Giga-Port project during which we completed SURFnet5.   This was five years after the end of that program and the effect on the economy of infrastructure components like

GigaPort had become very visible. This report appeared a few months before we had to defend the SURFnet7 project. So we ended up receiving a very positive endorsement for the GigaPort3/SURFnet7 project, both from the Netherlands Bureau for Economic Policy Analysis and the Commissie van Wijzen. The Commissie van Wijzen recognised at an early stage the importance of the integration of our different ICT Infrastructure components… including the e-science components. So this provided momentum for further governmental support.

Liebrand: One of the outcomes that we never expected was the pre-eminence of the Amsterdam Internet Exchange, the AMS-IX. AMS-IX and NetherLight are both examples of unanticipated payoffs from the SURFnet projects.

COOK Report: Did AMS-IX (along with all the other basic infrastructure) help Dirk van der Woude and Herman Wagter get the Amsterdam city fiber build off the ground including getting EU approval?

Liebrand: Yes. I fully agree with your conclusion. The technology position of Amsterdam is directly due to the investment in the networking infrastructure.

De Laat: And because of this, (the grid infrastructure and the knowledge we have,) EGI.org (the European Grid initiative) has decided Amsterdam is the place to be.

Liebrand: It is difficult to calculate the exact payback. But if you analyze it the way utilities calculate return on investment, then I would say there is a factor of four or five.

Neggers: The point, of course, is the spill over effect. The entity that does the investment will not see the returns coming back directly to it. They materialize elsewhere in the economy. Because this is infrastructure, it affects all sectors. This is why government must play a part as the investing agency that best crosses sectors.

SURFnet7 Builds on Nortel Common Photonic Layer, Next Generation Ethernet and E-science Middleware

COOK Report: I understand that SURFnet7 is funded, but what exactly is it?

Neggers: SURFnet7 builds on the foundation of the hybrid network of SURFnet6. The Nortel common photonic layer is still state of the art. Therefore, we will build on that.

COOK Report: How will Nortel's bankruptcy affect support for that technology?

Neggers: We are not worried. The moment it was announced, big competitors knocked on our door and said, "Don't worry. We will be able to maintain your network instantly should Nortel ever actually fail." There is such a large installed base that there is a solid market for maintaining this equipment. We believe there is little risk to the network, and we believe the technology is so valuable it will emerge intact from any take over. [**Note**: *Ciena recently announced that it has signed agreements to purchase optical and Carrier Ethernet assets of Nortel's Metro Ethernet Networks (MEN) business. Continuation of Nortel's photonic product line seems assured.*]

SURFnet6 was composed of five optical rings connected by optical electrical optical converters. We are removing these converters and will replace them by optical switches. This will make all of the Netherlands a single optical network — one no longer dependent on any kind of regeneration. This will permit instantaneous point-to-point connections anywhere in the Netherlands. Today a circuit switched light path from the

Year Work Package	2009	2010	2011	2012	2013
1. Photonics (PHO)	• Flexible photonics introduced into the network	• 40G waves introduced • 100G tested • New Planning tool • Study new technologies	• New photonics features introduced	• Technology Assessment • Architecture Study into next gen photonics for SURFnet8	
2. Next Generation Ethernet (NGE)	• Technology Assessment • Architecture study	• Procurement process SURFnet7 • Proof of Concept tests and preparations for new service provisioning	• SURFnet7 NGE services up and running	• New NGE features introduced	• Evolution into next steps • Design for SURFnet8
3. Enabling Dynamic Services (EDS)	• Multi-domain dynamic lightpaths Proof of Concept • Architecture Study	• Federated and Multi-domain lightpath Services introduced • Integrated resource services, studies and Proof of Concepts	• Proof of Concept for e-Science Service	• e-Science Services available	
4. Nether-Light and global con-nectivity (INT)	• NetherLight upgraded to scalable new platform	• NGE services available at NetherLight • Additional 10G intercon-tinental lambdas available	• NetherLight ready as "dynamic GLIF Open Lightpath Exchange	• New NGE features introduced in NetherLight • 40G intercontinental lambdas available	
5. Mobility and Fixed-Wireless (MOB)		• Technology and legal Assessments • Operator & Vendor scans	• Wireless application Living lab established with operators	• Wireless technology and mobile application Pilots with operators	• Hetero-geneous service offering
6. IP Innova-tion (FIP)	• Core router upgrade	• IPv6 support for connected institutions • Testing with 100G interface • Network Operations Center renewed	• 40G or 100G core trunks introduced	• Architecture Study future of IP routed networks	
7. User Par-ticipation & Knowledge Dissemina-tion (DIS)	• EYR Lightpath contest • Contact with research-projects	• EYR 3 preparations • 3-5 new research-projects • Knowledge domains • Events	• EYR 3 contest • 3-5 new research-projects • Knowledge domains • Events	• EYR 4 prep. • 3-5 new research-projects • Knowledge domains • Events	• EYR 4 contest • GP3 closing event • Knowledge domains

GigaPort3's main deliverables over the years, highlighting 2010.

south of the Netherlands to the north must pass through three optical rings, with OEO regeneration at each ring transition. SURFnet7 means this is no longer needed. We will also optimize the topology to make it even easier to go from a to b anywhere in the Netherlands.

We will move these lightpaths to a more dynamic user controlled interface. Using a GUI web interface, you will be able to decide where in the Netherlands you connect your light path. And we hope to extend this globally via the GLIF.

This is the Layer One optical part. On top of this we are planning to introduce a next generation Ethernet service as Layer Two. We call it "next generation Ethernet" because details of the new Ethernet standards are still in flux.

COOK Report: What are the parameters that will define the next generation of Ethernet?

Neggers: Actually it is fairly simple. It must be a scalable service. Ethernet today, in the sense of circuits, is not scalable. You have a limited number of VLANs, and globally you will not be able to build any really useful network out of it. Even the Netherlands is too big for this if you want to use it all over

– that is, in large scale. The major purpose of next generation Ethernet is achieving this scalability.

Local networks are Ethernet and will stay Ethernet. Ethernet connects everything. That means we will have a very user-friendly network layer at layer two. Of course Ethernet is also much cheaper than SONET/SDH equipment. Currently the optical, electrical optical conversion in the photonic layer is SONET/SDH based. SURFnet7 hopes to get rid of SONET/SDH entirely so that in the end it is just photonics and Ethernet if we are successful and the equipment market supplies it.

Liebrand: This will take another four years.

Neggers: Yes.

De Laat: This is also a research question, because SONET/SDH uses TDM and gives you deterministic behavior. In Ethernet you depend on certain protocols to separate traffic and the leaky bucket mechanism to shape and police each flow depending on the QoS parameters you set. We did experiments with Nortel last year. Our experiments showed video traffic could be run in such a network if protected from high amounts of "noise" traffic. This seems to be a significant step in confirming

where we want to go with next generation Ethernet. But more studies are necessary.

COOK Report: Will advanced networks like SURFnet7 leverage enable reducing electricity use and CO2 emissions by avoiding optical-electrical –optical conversion?

Liebrand: Yes. In fact, Bill St Arnaud is likely working on metrics about this. He wants everyone to measure their footprint and try to improve it in the next couple of years.

[**Editor's Note**: On December 8, 2009 Bill St.Arnaud replied to my query: No hard data yet. But the savings will be relatively small in going from OEO [Optical to Electrical to Optical] to OOO [all optical] with no electrical regeneration. Routers are the huge energy hogs and a OOO network significantly reduces number of distributed routers and allows you to locate core routers at a site that uses renewable power. SURFnet already realized this reduction of backbone routers in 2005 in the hybrid SURFnet6 network.]

De Laat: Here are some rough metrics. Each Fiber takes 80 colors and for each color you need a separate laser taking between 25 and 35 Watts. A complete rack would take 80 times 30 or 2400 Watts. If you use higher bandwidth you need

some more Watts. The state of the art photonics is such that you can reach much further distances nowadays without regeneration and have much better ways of correcting signal distortions caused by intermediate fibers and photonic devices. These new technologies of the last 3 years are now coming into products that make it possible to use wavelength selective switches [WSS]. Such a switch contains a small copper box that takes 9 fibers as inputs. The box is capable of selecting colors out of any fiber and inserting the mix of colors into an output fiber. And this copper box takes not more than 10 watts to operate.

Instead of 70 to 100 watts per lambda, a switch containing WSS'es can switch a few 100 lambas while using only 100 Watts for the photonic devices and controller. With this box you don't have to have a receiving and sending transponder for each color pair. You eventually need a transponder, but if you have the correct one, you can position them at the edge of the network.

Neggers: But getting back to the SURFnet7 plans – we will have higher speeds of course. One hundred gigabit/s will become the norm. We have already demonstrated 100 gigabit/s from Amsterdam to Hamburg

and back a distance of more than 1244 kilometers without regeneration. From Amsterdam to Geneva we have our own dark fiber that we will light with 40 gigabit/s lambdas from day one with no regeneration the whole way from Amsterdam to Geneva. You see therefore why we are not yet ready to sacrifice the Nortel photonic network.

De Laat: You should also note the impact of this on European networking where it becomes possible to cross national boundaries photonically instead electrically.

The point is that you are now able to extend the model of photonic communication to all countries. You can reach 1000 kilometers the distance from Amsterdam to Geneva which is about half of Europe without OEO conversion. You can also build Lambda Exchanges.

COOK Report: NetherLights? [See page 9 above.]

De Laat: Lower layer NetherLights.

cross national boundaries photonically instead electrically

Liebrand: But let's stay with

the technology for a while. It's good to notice that there are some policy problems that may impact on what the technology can do.

COOK Report: I conclude that if you can do this, then you have an infrastructure to which you can connect people in whatever way and by whatever means. Then you have an infrastructure that renders what belongs currently to the incumbents absolutely meaningless?

Neggers: Yes. But let's go in that direction more slowly. First of all I want to make it simple to connect and this is not the case today. Therefore an element of the SURFnet7 plan is called Enabling Dynamic Services where we want to make it simple for researchers to not just use the network but also to use resources available via the network in a very simple coherent way.

In short I am talking about the integration aspect where we have the data, the supercomputing and the e-science middleware to glue this all together. This middleware layer is also a main part of SURFnet7 development.

You have better and much faster photonics. Then you have Ethernet that makes it easier to connect and switch. Then you have the middleware layer that makes it easy

to bring in all kinds of re-sources and allow research-ers and people to collaborate on shared problems.

Liebrand: It is useful to see this in context.   This is a world class network in a small country, the Nether-lands, 200 by 300 kilometers and the moment you pass the borders you have also some political problems. It's desirable to notice that there are some policy problems that may impair the full po-tential of technology.

Neggers: Gordon is well aware that there are people in Europe and elsewhere who will want to slow you down if you go too fast. These are some of the headwinds we face in our march to the fu-ture.

Science is a collaborative effort that crosses boundaries.

COOK Report: Let's assume you do this and it gets the expected and hoped for re-sults for the Dutch economy, because the infrastructure means being able to do things in the Netherlands that can't be done elsewhere. If it makes you bloom, thrive and grow, won't Germany, France and the UK have to take no-tice and begin to understand

they will be left out until they slip the "incumbent boat an-chors" off the ankles of their economies?

Neggers: Right.

Liebrand: No, I am sorry. It is more difficult. It is like the prisoner's dilemma. It is in our own interest not to succeed too well, but to share the profit with the oth-ers,   and to collaborate to-gether.

COOK Report: Collaborate with your sister research networks in other European countries?

Liebrand: Yes.

De Laat: Science is a col-laborative effort that crosses boundaries.

GigaPort3 – E-science to E-life

Neggers: Referring to the "netness issue" that we dis-cussed earlier at the Portland show with its author Sheldon Renan, one of the new ele-ments in the Gigaport3 pro-ject is to have full mobility and a seamless integration of next generation wireless with our fixed infrastructure. This is a longer-term issue be-cause the frequencies have not yet been auctioned. But we want to be there and try to create a private network for our research and educa-

tion community with capa-bilities ahead of what will be-come commercially in the end.

COOK Report: When you talk about GigaPort3 it re-minds me of Nico Baken's concerns. Where does Baken fit into all this?

Liebrand: Nico Baken par-ticipated in the SURF founda-tion as a member of the Sci-entific Council. His schedule is such that he can no longer afford to spend a large amount of time with SURF, but he is well aligned with our goals.   As far as KPN goes, there is a big difference between KPN's research and development and KPN's commercial part.

Neggers: KPN does very good R&D, but from our point of view they are slower to bring the results of that R&D into commercial networks than we would like to see. We think KPN will be very interested in this wired wire-less interworking. Why? Be-cause we at SURFnet are not proposing to blanket the Netherlands with new radios. But we want to make use of the investments of public op-erators, and develop coop-eration between public and private investment.   There are enormous cooperative possibilities.

De Laat: Certainly true for the sensor grid.   The pro-

The Dutch Exhibit at portland Supercomputing 09. Unfortunately shot after closing. On Cees de Laat's website http://sc09.delaat.net/ you can find all the details of the SC09 Dutch booth and its demo's.

grammable controllable network that I showed you last year is part of the strategy to be able to influence and control the infrastructure. Last year we demonstrated programmable sensor networks for monitoring dikes (see: http://sc08.delaat.net/). With wireless many more sensor networks are imaginable.

applications and networks need to be aware of each other

The network itself is not sufficient anymore. It needs to

reach out to the applications and to the environment. The applications and the networks need to be aware of each other. A lot of the work that we do in describing infrastructure and applications

attributes in ways that the semantic web can be applied to the whole as an interlinked fabric. Making the relationship between the networks and everything connected to them is very important.

Some tile displays at the Holland booth in Portland. These are being fed from Sara computers to the booth on the show floor in Portland. The top tiles are display of changing ocean temperatures with the continents in white. The canal and birds and four times HD being streamed more than five thousand miles.

COOK Report: When the semantic web convergence with the service web is complete then these translations will become almost automatic?

De Laat: Yes - for example I can show that there is a lot of semantics research going on for applications, e.g. in biology DNA sequencing. There is also a layer of semantics in media content. And if we can link these semantics to infrastructure you get location aware content and location aware applications. And then you can really think about how to optimize the work that you want to do.

COOK Report: Is what you are doing in e-science intended to transition into what you might call e-life? Will the tools and infrastructures that you are building now be transitioned to support the greater virtualization of everyday infrastructures?

Neggers: Of course.

COOK Report: Isn't that a very deep and important part of what you are doing?

Neggers: This is exactly why the government and in particular the ministry of economic affairs is supporting us with additional funds and are not just leaving us to the support of the ministry of re-

search and education. It is the spill over effect to society as a whole that justifies this broader support.

De Laat: Definitely the next step after virtualization and new methods of e-science is the virtualization and new methods of e-life in the community. It will take another five to ten years to get there but, as I also tried to point out yesterday, life as is currently lived on earth is not sustainable for ten billion people.

Neggers: As a very straightforward example, Philips is slowly dropping out of consumer electronics and moving into health care. And for them health care means a connected person. Wherever he is and that is why you need the seamless interworking of wireless and wired to make that all happen. For them this is their core business. If the infrastructure is not there, new tools to support better health care will not be there either.

De Laat: "Our monitoring show that you are due for a heart attack in five minutes, please come to the hospital."

COOK Report: Sheldon Renan told me that he has had a lot of discussions about this with Jacob Bardram who is the head of pervasive computing for healthcare in Denmark. And we also have

Sheldon's example of the bathtub that catches the person when he falls.

Neggers: It is this kind of lateral transectoral thinking that we want to trigger with Enabling Dynamic Services. It is the horizontal infrastructure that glues all this together.

COOK Report: Is the going wireless component part of SURFnet7 or will it be a part of SURFnet8?

Neggers: In the GigaPort3 plan we will investigate this together with industry. SURFnet8 likely will bring the integration of what we learn with GigaPort3 into the network fabric that we are building. We will ask government again for support to make this integration happen in SURFnet8.

COOK Report: And a significant reason for on going support is to keep the Netherlands competitive with the larger countries?

Liebrand: Yes. Absolutely.

COOK Report: This is the same path that the Danish government has followed very productively.

Neggers: Did you know that wi-fi is a Dutch invention?

De Laat: In fact it was invented two kilometers from

my house and involved my neighbor as I found out when I was digging into application of wi-fi in education 14 years ago.

Neggers: Lucent at that time had three labs in the Netherlands.

The Sorcerer's Apprentice Race Between the Creation of More Bits and the Means of Storing Them

COOK Report: Let's explore e-life more. I remember that Sheldon Renan recounting a talk John Seeley Brown gave in 1994 where he explained that typically you cannot sustain productivity increases beyond 10-12% per annum in the world of atoms. But in the digital world - if you combine the annual improvements in productivity of both hardware and software - you may be able to achieve and sustain productivity increases of 100% per annum. The way John expressed it, although Moore's Law is focused on hardware (semiconductors), improvements in software appear to be more significant.

Software is much faster than hardware. New software running on old hardware will outperform old software running on new hardware. Experi-

ments at Xerox PARC concluded that anything you could do digitally would be much more successful than things done in analogue fashion.

Both Sheldon and JSB are saying that Moore's Law is no longer the most important source of new opportunity. JSB said software improvements were more significant than hardware improvements. And Sheldon is saying that connectivity improvements are more significant than hardware improvements. Sheldon also points out that as you reach a situation in the real world where you're running out of resources, you are forced into the digital world.   That is your only hope apart from a Malthusian reduction in population.

connectivity improvements are more significant than hardware improvements

[**Editor's Note:** –in the spring I will publish an in depth exploration of netness (also known as Renan's law) with interviews done in Portland, November 17th and 19]

De Laat: The same holds true with e-science as it becomes a global and multidisciplinary effort as you need

the generic tools that allow the data of different disciplines to be combined to enable the next steps of further progress.

If you can't do that and your competitors can, then you loose.

Neggers: But there is also another challenge for the world of bits. In the data environment, the creation of bits grows like mushrooms because bits are so easy to create that like the sorcerer's apprentice it becomes difficult to put any limits on their creation. Storage of them takes a lot of resources. The question becomes whether the increase in storage capacity is outpacing the production of bits.

COOK Report: And doesn't the question become one of whether the environmental impact of electricity and cooling needed is sustainable?

De Laat: There is another issue. While you may eventually get the capacity of the library of congress into a single disk drive, you will find that scientists in massive data producing experiments as in LHC tune the amount of data they keep. As the capacity of available storage increases you will find that they measure way more variables because keeping more data allows them to do more sophisticated analysis.

The film industry would like to store everything that comes out of Hollywood for 200 years. They are moving from 4k to 8k and now everything is done digitally.   They also want to store the workflow, everything involved in making each movie. Storage needed per movie will be between one and three petabytes.   They produce 600 movies per year that may make it to academy award (Oscar) consideration process. They want to keep all 600 because they have no way of knowing today what will be the cult movie 50 years from now.

If you look at these data requirements, the requirements of the Large Hadron Collider are mere noise.   You must look at the data evolution curve. People experience 2k and then 4k and then they want 8k.

The general lesson for this discussion is that as the infrastructure grows faster and more capable, scientists will find new problems that will continue to stress the infrastructure no matter how fast it grows.

COOK Report: I sense that there is an evolutionary process for conceptualizing what you are doing.   Would you elaborate further?

Neggers: Well as Wim says the government gives you an

innovation budget but at the same time the users are committed to pay for the use of the new infrastructure in the end and, therefore, they are also deeply involved in the process of developing it. We are not a laboratory that has a few techies on board and tries to do nice technical things and then hope that they will get used. To the contrary we operate with constant feedback from our user community as represented by Wim Liebrand. SURF is managed by the users and not by the government. The government trusts the users to integrate the infrastructure and not the ministry. The users are in charge of their work.

Liebrand: SURF embraces two communities: the techni-

cal and research community and the end user community. We are trying to interconnect them.

Neggers: Some people say that because our citizens have been forced to work together to protect themselves from the rivers – actually floods from the river are more serious a threat than the sea – the ability to have a common responsibility for the greater welfare is ingrained.

> ## SURF is managed by the users and not by the government.

COOK Report: A source of

the pragmatism?

Liebrand: It is true we have to cooperate and be competitive. But also can't forget the role of the pirates. In other words, if you excel at building a world class infrastructure, others will become jealous. And when jealousy rules, collaboration fails. By some degree of trial and error we have discovered you need to excel in certain fields, but also be willing to share with your competitors.

COOK Report: And if you do not do this?

Liebrand: In the long run you will loose the game.   It is again basically the prisoner's dilemma.

Neggers: This is why we are so active in global collaboration. Because it would be useless to have a nice network that connects everyone inside the Netherlands and nowhere else.

COOK Report: What you are doing is so powerful. What I find ironic is that the power is not well recognized, nor is it understood.   People tend to think in terms of their own specialties. No one does a lot of connecting of the dots. But as you saw, when I introduced you to Sheldon Renan, his focus on the interconnectedness of everything may begin to help people see the larger value.

Left to Right: Sheldon Renan, Kees Neggers and Cees de Laat at the Dutch Booth. Sheldon is giving an informal exposition of Renan's Law - Everything wants to be connected to the "flying" Dutchmen. All parties agree that these are important ideas -detailed exposition coming soon in the *COOK Report*.

IV. The Ascent of e-Science

Promise of the 4th Paradigm

Introduction e-Science as the Fourth Paradigm

Editor's Note: Twenty years ago I was privileged to be immersed for almost three years in Computational Science at the John von Neumann National Supercomputer Center in Princeton. I understand the first three: empirical, theoretical and computational. The fourth is fall out from two more decades of Moore's law – data intensive in that the computational simulations of twenty years ago are now extended in hitherto unimagined ways yielding experiments that produce huge amounts of complex data.

The following paragraphs are taken from a new book The Fourth Paradigm dedicated to the memory of Jim Gray. http://research.microsoft.com/fourthparadigm.

The introductory chapter of the *Fourth Paradigm* is "Jim Gray on e-science: A Transformed Scientific Method" based on the transcript of a talk given by Jim Gray to the NRC-CSTB1 in Mountain View, CA, on January 11, 2007. Early on Gray says, "Look, computational science

is a third leg." Originally, there was just experimental science, and then there was theoretical science, with Kepler's Laws, Newton's Laws of Motion, Maxwell's equations, and so on. Then, for many


problems, the theoretical models grew too complicated to solve analytically, and people had to start simulating. These simulations have carried us through much of the last half of the last millennium. At this point, these simulations are generating a whole lot of data, along with a huge increase in data from the experimental sciences. People now do not actually look through telescopes. Instead, they are "looking"

through large-scale, complex instruments which relay data to data centers, and only then do they look at the in- formation on their computers."

'The world of science has changed, and there is no question about this. The new model is for the data to be captured by instruments or generated by simulations before being processed by software and for the resulting information or knowledge to be stored in computers. Scientists only get to look at their data fairly late in this pipeline. The techniques and technologies for such data-intensive science are so different that it is worth distinguishing data-

intensive science from computational science as a new, fourth paradigm for scientific exploration.'
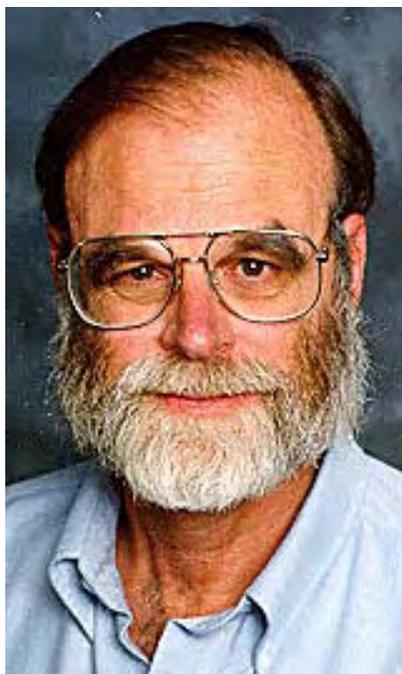
Jim Gray 1944-2007
http://en.wikipedia.org/wiki/Jim_Gray_%28computer_scientist%29

"We are seeing the evolution of two branches of every discipline, as shown in the next slide [Figure 2]. If you look at ecology, there is now both computational ecology, which is to do with simulating ecologies, and eco-informatics, which is to do with collecting and analyzing ecological information. Similarly, there is bioinformatics, which collects and analyzes information from many different experiments, and there is computational biology, which simulates how biological systems work and the metabolic pathways or

the behavior of a cell or the way a protein is built.'

'This is similar to Jeannette Wing's idea of "computational thinking," in which computer science techniques and technologies are applied to different disciplines.'

'The goal for many scientists is to codify their information so that they can exchange it with other scientists. Why do they need to codify their information? Because if I put some information in my computer, the only way you are going to be able to understand that information is if your program can understand the information. This means that the information has to be represented in an algorithmic way. In order to do this, you need a standard representation for what a gene is or what a galaxy is or

what a temperature measurement is.'

Experimental Budgets Are ¼ To ½ Software

'I have been hanging out with astronomers for about the last 10 years, and I get to go to some of their base stations. One of the stunning things for me is that I look at their telescopes and it is just incredible. It is basically 15 to 20 million dollars worth of capital equipment, with about 20 to 50 people operating the instrument. But then you get to appreciate that there are literally thousands of people writing code to deal with the information generated by this instrument and that millions of lines of code are needed to analyze all this information. In fact, the software cost dominates the capital expenditure! This is true at the
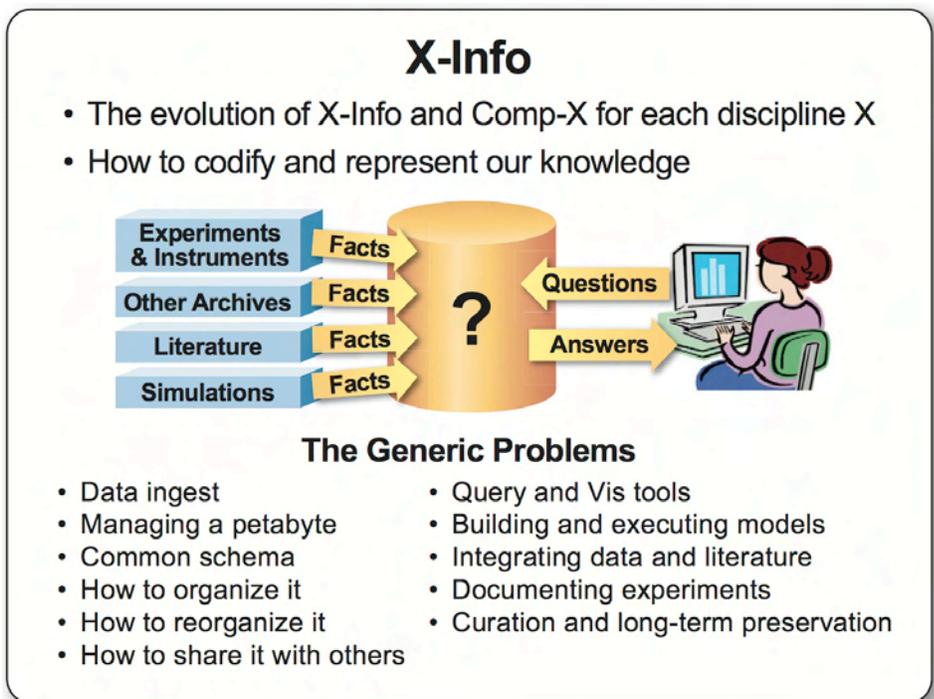
Sloan Digital Sky Survey (SDSS), and it is going to continue to be true for larger-scale sky surveys, and in fact for many large-scale experiments. I am not sure that this dominant software cost is true for the particle physics community and their Large Hadron Collider (LHC) machine, but it is certainly true for the LHC experiments.'

'Even in the "small data" sciences, you see people collecting information and then having to put a lot more energy into the analysis of the information than they have done in getting the information in the first place. The software is typically very idiosyncratic since there are very few generic tools that the bench scientist has for collecting and analyzing and processing the data. This is something that we computer scientists could help fix by building generic tools for the scientists."

"I have a list of items for policymakers like CSTB. The first one is basically to foster both building tools and supporting them ..." *Fourth Paradigm*, pages xviii – xx

And Then Jim Gray Vanished . . .

Despite human desires and especially those of scientists to explain and know everything -- just when you think you know, and are sitting astride the world with the progress you are making, fate sometimes rudely interrupts. So it happened with Jim Gray.   Suddenly and tragically and mysteriously.

We read "During a short solo sailing trip to the Farallon Islands near San Francisco to scatter his mother's ashes, his 40-foot yacht, Tenacious, was reported missing on Sunday, January 28, 2007. The Coast Guard searched for four days using a C-130 plane, helicopters, and patrol boats but found no sign of the vessel.

Gray's boat was equipped with an automatically deployable EPIRB (Emergency Position-Indicating Radio Beacon), which should have deployed and begun transmitting the instant his vessel sank. The area around the Farallon Islands where Gray was sailing is also well north of the East-West ship channel used by freighters entering and leaving San Francisco Bay. The weather was clear that day and no ships reported striking his boat, nor were any distress radio transmissions reported.

On February 1, 2007, the DigitalGlobe satellite did a scan of the area, generating thousands of images. The images were posted to Amazon Mechanical Turk in order to distribute the work of searching through them, in hopes of spotting his boat.

On February 16, 2007, the Friends of Jim Gray Group suspended their search, but continue to follow any important leads. The family ended its search May 31, 2007. The massive high-tech effort did not reveal any new clues."

http://en.wikipedia.org/wiki/Jim_Gray_%28computer_scientist%29

Just when one would be certain that the tools of Gray and his colleagues would have found evidence. . . they did not.   Mystery seems to have had the last word.

V. Making E-science Work: The Middleware Solution

An Interview with Bob Hertzberger

L. O. (Bob) Hertzberger is the Director of the Virtual Laboratory e-Science project, better known as VL-e, an essential contributor to the National ICT Research Infrastructure for The Netherlands.

Hertzberger leads the critical VL-e initiative to develop the functional model that, together with the methodology of multi-disciplinary collaboration, is shaping a new generation of advanced e-science software. Hertzberger's goal is to provide generic middleware that is easily adaptable to many disciplines in order to make the technologies of "The Fourth Paradigm" affordable and accessible to many more researchers.

This is expected to provide a critical component in the Dutch effort to create a national knowledge infrastructure. But as can be seen in the interview with an American medical researcher in Section IX, the new model and resulting benefits are having a global impact.

The COOK Report met with Bob Hertzberger at the Portland Convention Center on November 17, 2009. The edi-tor asked him to provide some context around his new approach to developing middleware at the VL-e.

Hertzberger: I would prefer to modify the Science Paradigms slide in the previous section to observe that till around 1930-1940 theoretical science and empirical science had developed to such an extent that complex theories were laid out and complex experiments were conducted in order to understand the phenomena of *empirical science*. Above all, this was the case in physics. The steam engine was developed based on these insights. This, in turn, laid the foundation for the industrial revolution and railroads. Later the combustion engine emerged that helped to make the automobile and the airplane possible.

...e-science must be able to combine data-centric and computational models.

Bob Hertzberger at the Portland Convention Center on November 17, 2009

By that time the interplay between theory and experiment necessary to understand empirical phenomena had progressed in such a way that theorists developed complex models. Experimentalists carried out complex experiments partly to verify the theorist's models, partly to search for new phenomena.

One of the best examples is the Bohr model for understanding atoms and molecules that constitute the various elements like iron, copper, gold etc. (For lay persons: the Bohr model says the each atom or molecule of an element is composed of a nucleus with electrons circling around it like planets around the sun.)   Before the Bohr model various other models had been developed. But when these models were tested with the combination of theoretical work and experimental evidence, the results demonstrated they did not behave as predicted by the empirical observations that most elements are stable.

To carry out such experiments, scientists have used instruments such as the microscope to study all kinds of small objects like the cell, or the telescope to study the planets or stars or galaxies. But in the mid-forties of the previous century there emerged a new instrument which was called the com-

puter. The computer has turned out the most flexible instrument of all.

In the beginning it was rather clumsy and very limited. But as its speed and storage capacity increased, the computer became capable of making far more complex calculations than a human being. It could also store all kinds of digital information more effectively.

The computer made it possible to experiment with models what we call simulations. Simulations could represent a certain phenomena   — for instance the impact of CO_2 absorption in the atmosphere. It could be executed by the computer with a parameter of interest as defined by the scientist (for example the impact of burning fossil fuels that then release CO_2 into the atmosphere) set to a certain value. After that we could change the parameter, run the model again and observe the effect of our change. These types of models have been designed for a large number of sciences including physics, chemistry, and pharmacy. But modeling and simulation could also be used to make economics forecasts or study road traffic or social behavior, etc.

The computer could also be used in "real world" experiments to gather the data in a digital form from instruments

like a microscope or a telescope. That data could then be subjected to complex processing, and the results analyzed by humans to produce findings that would give more insight into, for example, the workings of a cell or the shape of a galaxy. Today, the most powerful application of computer instrumentation occurs when we combine the results from simulations with those of real experiments. Doing so will yield better simulation models.

The computer has turned out to be the most flexible scientific instrument of all.

Twenty five years ago we began to connect computers together into networks. These connections resulted in the emergence of the Internet and Web. Over the past decade copper wires composing these connections are being replaced by optical fibers making it possible to communicate by means of high bandwidth at virtually the speed of light. This has resulted in what we call "globalization", where data and information about economics, society and business can be communicated and shared globally on an unprecedented scale.

The Promise of "The Fourth Paradigm"

Data and information can now be seen as the core component of science, as the computer penetrates ever deeper into the profession of the scientist. Accessing, handling, sharing and combining the resulting data and information will become more and more important. Optical networks make these processes easier and faster.

The consequence of the emergence of computers, especially when connected via fast networks, is that a new form of science is emerging which is called "e-science". It enables the practice of science at a global scale by means of communication over the internet. For this form of science to be successful, it is necessary that data and information produced with public funds stay public and that access (for instance via public databases) be supported and facilitated. It could well be that such an e-science computer and network infrastructure will turn out to be the most flexible instrument ever designed. We are working in the Netherlands to realize this potential.

COOK Report: Your context for Jim Grays' last lecture is also the foundation for your own work. Optical networks

are the basic foundation for the connection of scientists to major instruments and sets of data gathering tools. Such networks must be global in scope, and the national agencies that support them must figure out how to cost effectively allocate funds for the network, and the instruments that the network ties together, as well as the server clusters that record the data generated.

…we develop a systematic approach to e-science

And as Cees de Laat pointed out earlier, the cost of the software – which makes it possible to do something useful with all the hardware -- can exceed the cost of the hardware. Software has too often been an afterthought. The Dutch Roadmap for SURF gives the greatest emphasis to the need for creation of this e-science middleware stack on which you are working. What is the story behind your new software?

Hertzberger: In the Virtual Laboratory for e-science project we have tried to bridge the gap between a state of the art network and being able to use such infrastructure for doing new science. We have begun by develop-

ing a systematic approach to e-science as a whole, in order to provide a framework for creation of the various software tools necessary to realize an e-science environment.

Via this model we intend both to stimulate the creation of communities of best practice and to enable multi-disciplinary collaboration. We also believe our model helps to implement the hardware and software environment necessary to realize e-science.

Our model delineates dual forces. On the technology side are the computer, storage and networking environment, including all software needed to harness that environment like Web, Grid and Cloud. We call this "technology push".

On the application side, we describe the instrumentation environment needed to realize experiments including simulation experiments. Our application side includes all the software tools needed to acquire, process and analyze the data so that the information and knowledge necessary for the scientific understanding is produced. We call this "application pull".

The environment needed to accomplish our work should not only contain methods and tools that we have produced ourselves but it should also be capable of incorporating

tools made by other scientists. That is why, besides realizing the application pull model and researching new methods and tools ourselves we also have focused on methods and tools needed to incorporate existing e-science tools. For instance we want to be able to use the Taverna workflow tool [See *4th Paradigm* p. 137] in our environment. We have no intention of wasting resources to re-invent the wheel.

Information and data become the core business of e-science…

While we concentrate on the application pull of e-science, we remain dependent on changes in the technology push. Ideally we want to become as independent as pos-

sible of the technology push exerted by developments in Web, Grid, and Cloud computing. In practice, however, this is never completely possible. As an example, when Grid changed from a process into a service-oriented environment, we had to adapt to this service orientation. Similarly as Cloud computing becomes more mature and standardized we will have to adapt our system to meet the demands of this new technology push. These dynamics also hold for the Web, especially as the semantic Web matures and is incorporated into current Web standardization.
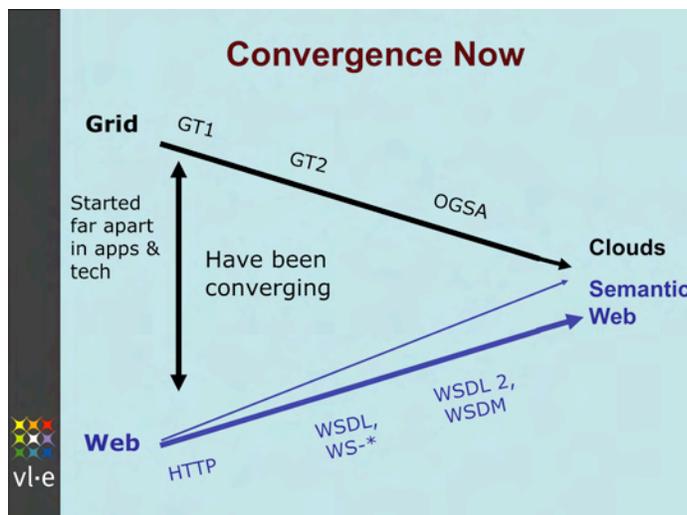
The Maturation of Our Application Tools

Hertzberger: The slide titled VL-E Project on the third page down represents the overall structure we are building to become, as much as possible, independent of the Cloud or Grid. It is coming into shape as a set of tools represented by a layered structure. The purpose is to enable researchers to have the most powerful and effective software possible to connect their workstations to the resources, whether they are databases, scientific instruments or high performance computing located in the global grid or cloud.

To help researchers who use clouds and grids, the field has created numerous tools such as Carole Goble's Taverna. Taverna is a very good example of an e-science tool. These tools are designed to be used by a large commu-

nity to capture their methods and process their data and information. Now each community tends to have its own sociology.

In life-sciences for instance the emergence of new instrumentation has made so-called omics-based experiments (transcriptomics, proteomics, etc.) possible. The omics technologies form the basis of a new form of (molecular) biology that cannot be realized without extensive use of computer technologies. Its consequence is an increase in complexity of the experimentation as well as data and information that arrives in greater quantity and complexity than ever before. This has had a tremendous impact on experimentation in biology, pharmacy and medicine. e-science as a data centric science helps in harnessing the (computer generated) complexity impacting the basic methodology by means of which domain science (in this case life science) is accomplished.

As a next step we have to demonstrate that our e-science model can be used to handle data from instruments as well as data generated by computational simulations. I am confident that will be possible. In e-science you need access to both data-centric and simulation models to help you understand scientific problems. An example would

be the study of system biology.

COOK Report: While the software could be open source, and most of the databases could be open, but the processes and results have not necessarily to be open?

Hertzberger: Yes. A pharmaceutical company may well use open tools but as long as all its critical knowledge and data are wrapped into a proprietary database under its own control, that company should feel secure.

Let me loop back to life sciences and the pharmaceutical industry with which I am quite familiar. Here you would expect that their databases about the clinical trials which are their core business will be kept closed. However, the entire infrastructure could use open software tools. Nevertheless, at the end, some of the databases that house their results, and the management of their data will need to be closed (protected) because they contain company sensitive information. As long as companies can protect critical data, information and knowledge, there is no reason why public and private research parties should not work together.

COOK Report: What possibility do you see that researchers will begin to do

their research in more open ways?   Wouldn't that research be more effective than a closed proprietary approach?

Hertzberger: In the Netherlands some plant breeding companies are very afraid of Monsanto because they patent a lot of seeds. Monsanto can do that because they have more powerful development facilities of which important parts are bio-informatics tools. But I have told a lot of people not to worry because if you build your own e-bioscience chain you can compete with them. Why?   If they keep their chains proprietary, they will have more software costs and less flexibility to use other people's tools. We have an example of a pharmaceutical company in the Netherlands who for decades, at considerable costs and not much success, has tried to set up a proprietary e-bioscience chain. It is now embraces the semi-open collaboration model I am advocating here.

An open environment can be far stronger, because many more scientists can contribute to the power of the chain, cross checking it and contributing. Last but not least, being part of an e-science environment means that you don't have to support all the tools yourself, especially when they are public domain.

So I agree with you that science will be a lot more productive when all these new e-science efforts are brought to fruition.   These ideas are similar to those of Tony Hey. He started the UK program in e-science. Now Tony is at Microsoft where he is working towards this end.

COOK Report: Doesn't Microsoft have a proprietary attitude? There must be real limits on how proprietary you can be and still be competitive in e-science.

Hertzberger: Precisely. That is the whole point. Openness via e-science makes you, as a scientist, more competitive. But let me offer another insight. That I have ended up in e-science at the end of my career is because I had a splendid career in high-energy physics. High-energy physics could never even exist without international collaboration. This is now coming up more and more in other fields of science as well.

…openess via e-science makes you as a scientist (or organization) more competitive.

Have lunch in the cafeteria at CERN, and you will meet all kinds of people who can

teach you more in a few days than you could learn from books in months. I am working in e-science now because I believe that you must develop tools that make this kind of learning, learning via electronic means, possible for people in all fields of science. Not just in high energy physics.   Personal contacts are essential in science, and those can be supported by e-science through tools for supporting collaboration and structuring multidisciplinary interactions.

Collaboration with industrial partners can be realized via public/private partnerships for pre-competitive research. This can cover the company's end domain core business (for instance, pharmaceutical, biological or medical research). But it could also be in the form of R&D necessary to enable an e-bioscience problem-solving environment. A well-designed e-science infrastructure will help everybody including those companies.

Here we are talking about doing science on a global level. Using an e-science environment you can build different scientific communities -- public, private or mixed.

COOK Report: My attempt at defining this platform would be to call it the set of infrastructure tools that becomes the medium through which he or she can choose

and shape their interaction with the database telescopes, other instruments and supercomputers in cooperative exploration with colleagues around the globe.

Hertzberger: That is correct. In the functional model we are using and that is presented on next page (people tease me by saying that it looks like a multi-stacked ocean-liner) the towers standing on the horizontal slabs represent, in fact, the different user communities. The total technology chain necessary to create the problem-solving environment a scientist needs for his work is made up of the software represented by the towers and all the underlying software and hardware layers. This software could reside both on the scientist's laptop and be located elsewhere as well — for example in Web, Grid or Cloud based servers. It should allow him to do his experiment and connect with fellow scientist when necessary. But it requires in-depth e-science research to find ways how to use it.

Let's assume we have a scientist working in the area of medical and cognitive applications, and furthermore that he wishes to collaborate with a colleague in the U.S. working in the same area. Now imagine that they do an experiment in cognitive brain research collaboratively. They

use the green silo as the source or the "basis" of their software for doing the experiment. By means of the same software they can communicate with the database located somewhere else in which the data of some of their previous experiments as well as the results of that experiments are stored. Together, they use the tools necessary to have access to their shared data in the database and have the tools – perhaps workflow tools - necessary to complete a new study. After finishing that study, they can store their findings, which they could choose to make public or to keep for themselves until they get a better understanding of the issues.

If they decide to make the database public, they can announce to the other researchers in cognitive science: the results from our study are here. Please have a look and comment. Or they could select others with different expertise and invite them to join their work.

COOK Report: Then you get the widespread incremental benefits that the Linux community provides from everyone who is interested in the process and actively thinking about how to improve it. But while no one is going to force them to make their data public, will the "evolutionary" process that occurs as this
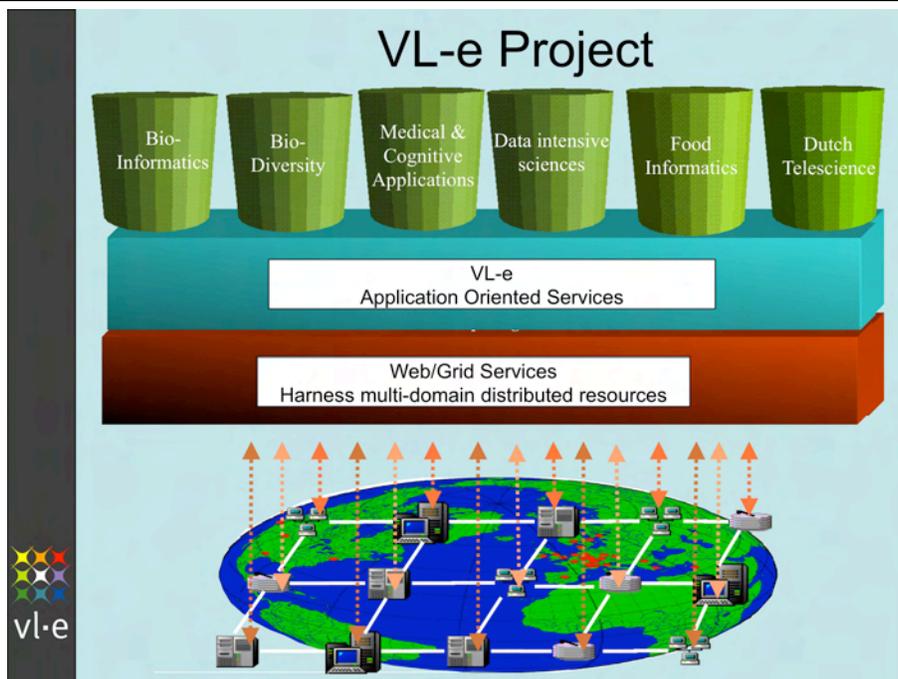
moves forward put increased pressure on people to choose open collaborative research?

Hertzberger: Certainly, but don't forget that researchers in a big pharmaceutical company like Merck will also have access to the same infrastructure without the need to publish things they want to keep private. In short, while the methods of doing research are changing drastically, the data and information becomes the core business of the scientist in both public and private organizations. This opens up new opportunities. But it demands in depth e-science to enable its application.

These research methods can benefit, for example, from the use of a workflow tool to describe the experiment that must be done. Such a tool is

very useful because it formalizes the structure of the experiment — and consequently make it more reproducible.

COOK Report: And semantic capture becomes an integrated part of the tools? Correct?

Hertzberger: Yes. From the beginning of the VL-e project we recognized that semantic modeling was of vital importance. Research involving the semantic web has made available knowledge representation mechanisms (for example OWL and RDF) that have gained general acceptance, together with a host of tools to manipulate semantic models. It is also clear that, given the speed with which data are gathered in e-science, in creating semantic meta-models, we must rely

on adaptive techniques. For this purpose we created the AIDA toolbox: http://www.adaptivedisclosure.org/aida

As an example: Semantic information is important in the context of Workflow composition, and this so even if automatic workflow composition remains a long term ambition. In order to incorporate a service in a workflow, one needs to know what the service can do and what kind of data it accepts and creates. This knowledge is essentially semantic information.

This work is similar to the example of Taverna by Carole Goble. She has built a workflow toolbox that she hopes the whole biological community will use and which will also capture semantic information. An other example is that semantic models will play a large role in the systems and network management in Grid and Cloud computing itself. Automatic configuration of systems and networks will be impossible without them.

COOK Report: I like these ideas very much. This past summer I read about something that would have been unthinkable in my youth. Clinical studies in the *New England Journal of Medicine* had been faked! It seems as though this would end those practices and doing that

would certainly be good.

Hertzberger: Correct. But doing that will be more difficult with e-science, because it makes science more open, controllable and reproducible.

COOK Report: Because it is reproducible, by definition it is verifiable?

Hertzberger: But now we could go a step further. It is the step advertised by Tony Hey and is beginning to take shape in some communities. In those communities you electronically publish your article together with the information needed to verify for others that the experiment will have the outcome that you claim. A publisher will tell you, if you want to publish an article you also have to publish the data and information on which that article is based.

In the future in biology publishing an article almost certainly will mean publishing the workflow and relevant data/information so that other people in effect can "replay" the experiment. However, we have not yet talked about the fact that e-science demands for a multidisciplinary way of doing science, which creates a question of who gets the reward. In the CERN example, new rules have been put in place to give everybody his/her reward in a multidisciplinary

collaboration. This is a big obstacle for "open" collaboration in those sciences that still depend on a more traditional publication reward system.

System Level Science

There is another consequence of e-science which might have an even more profound impact on science. This is the fact that e-science supports a system level approach towards science. In modern science and engineering, and in society, too, we are increasingly faced with complex problems which can not be studied separately. They can only be understood in the context of the full overall system they belong to. The study of this type of problems is often called as *system level science*. One of the best definitions for it is: *"the integration of diverse sources of knowledge about the constituent parts of a complex system with the goal of obtaining an understanding of the system's properties as a whole.",* from Ian Foster (Argonne National Laboratory) in the November 2006 issue of scientific journal IEEE Computer.

Some examples of *system level science* might be:
• study into the role of black matter for understanding the origin of the universe

- study of the fundamentals of matter in physics
- the role of understanding the functioning of the cell for system biology
- cohort studies in medicine (biobanking-based)
- environmental studies into the role of biodiversity in ecosystems
- studies of water and air pollution
- discovery of biomarkers or ligands for drug design in pharmacology
- modeling the behavior of continental shelves for earth quake prediction

Because of the potential of e-science to gather information on a global scale, it becomes possible to utilize all that data and/or information for doing system level research. However, it then becomes important to understand how that information can be integrated in order to add to the understanding of the complex problem under study.
In other words, how do you design system level models that are capable of integrating all that data and information?

COOK Report: You talk about the role of the contextual integrator. That seems absolutely critical. Wow is your approach to e-science different from that of others?

Hertzberger: In the Virtual Laboratory for e-Science (VL-e) we have been developing

a systematic approach to e-science as a whole. We do this to provide a framework for creation of the various software tools that are necessary to realize an e-science environment. We have developed a functional model, shown on the "VL-e Functional Model" slide below, with the intention of stimulating the creation of communities of best practice as well as multi-disciplinary collaboration.

…e-science is an enabler for system level science, an example of which is systems biology

Additionally the functional model supports the coher-

ence of developments in the project as well as the creation of modular software components. Our model must contain the problem-solving environments necessary for the different applications and application domains. In developing our model we made a distinction between those components (services) that can be realized as generic and those components which are more application-specific.

This makes it easier to modularize software components. It helps to create generic software where-ever possible and therefore encourage re-usability. And the goal – given the expense of software development - is to make software more capable, more powerful and more useable over a broader range of applications.

Focusing on modularization

VL-e Functional Model
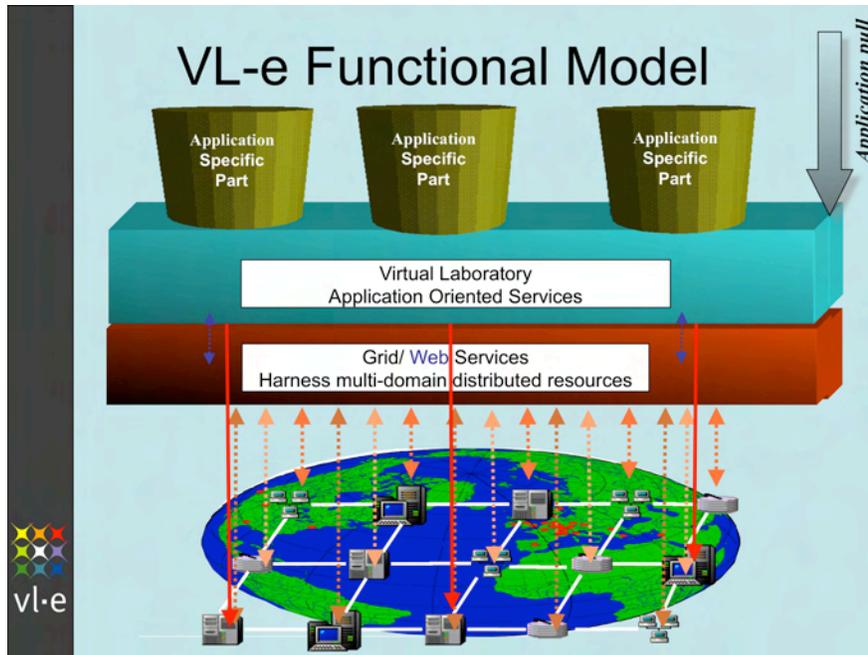
to lower cost also has the effect of influencing our other applications to become more generic. Workflow would be an example.

We have already explained that in realizing such a functional model, one has to observe two important developments: the push of the new technologies, and the changes taking place in those technologies. On the other side you have the pull – in other words, the demands from the various application domains like physics, medicine or pharmacy as they are placed into your total e-science system.

On the "application pull" side we minimally distinguish two different abstraction levels. One contains the application specific software. The other contains services that are generic for all e-science ap-

plication domains. All software has to be realized via the Software As a Service (SAAS) concept. It's important that all application services be designed as much as possible to be independent from changes taking place in the "technology push" environment. As an example, suppose the cloud is transformed into something else. Then we have to modify those modules that are designed for communication with the cloud environment.

COOK Report: So the generic modules sit on top of the evolving foundation? Can they maintain their viability independent of what goes on above or below them?
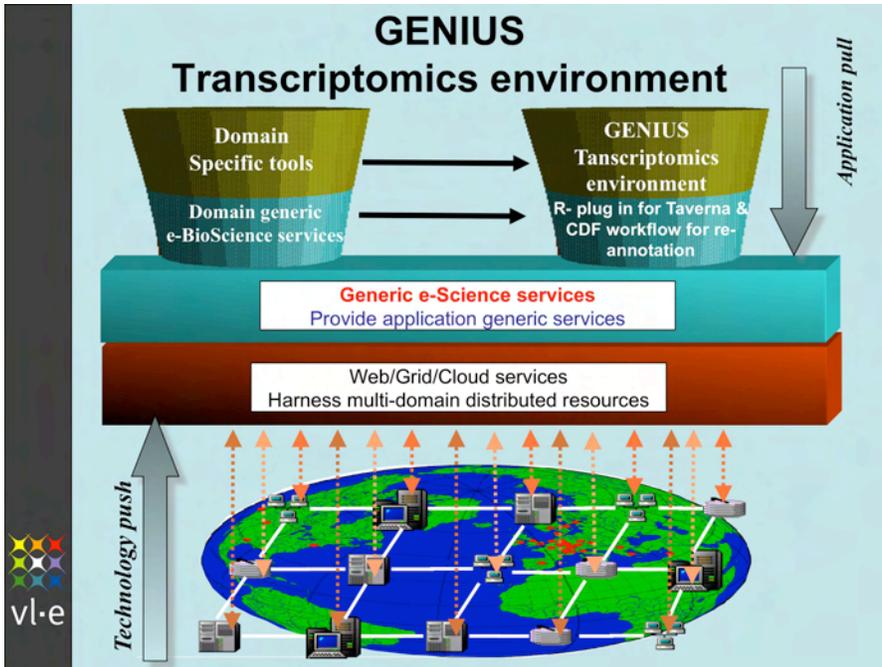
Hertzberger: The change in Grid from process to service orientation we talked about earlier required a major

change in the generic software as well as the application specific parts. To build an application, you apply as many generic services as possible so that the application-specific part is as minimal as possible. Then if you have to change your application, you only need to change a small part. You can leave the rest untouched. Although all services try to access the underlying Grid or Cloud services through a well-defined interface, that does not guarantee that when major changes in the underlying software take place, such changes do not have an impact on your total system.

COOK Report: Abstraction layers are determined by communities of best practice. As more of these communities emerge, will this require adapting the functional model?

Hertzberger. Yes, we must then refine the model. Probably we will add more abstraction layers. It took a lot of trial and error before we saw we had the current layers correct. For me the proof of the pudding is always in the eating. I needed to build the software to see whether the model was correct.

The model we discussed is presented in a general way with a clear distinction between generic and application
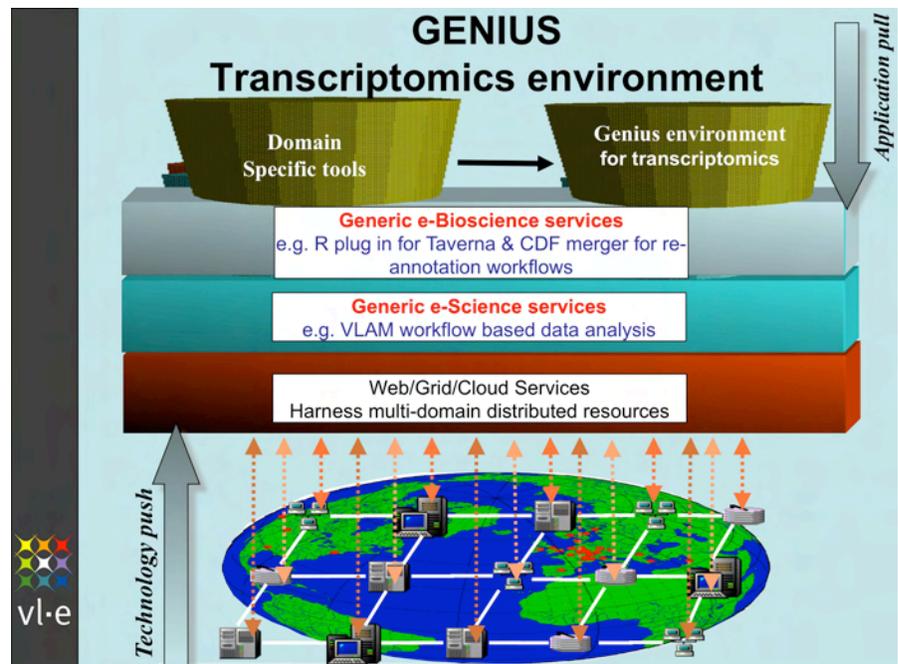
specific layers.   I wants to use it in this discussion to illustrate that because of the work we've done, we already believe that more generic layers are possible. Consider the work going on in VL-e currently to support one of the basic genomics technologies. These are studies based on micro-array instrumentation. They're called transcriptomic studies. This is a specific instance of the general process of how we build problem solving environments in VL-e. The process is similar whether it's for biodiversity or food or medical or cognitive research.

A Technology Chain for Micro-array Tanscriptomics

Hertzberger: In order for a scientist to do transcrip-

tomics research, i.e. genomewide expression profiling experiments, he needs to have a transcriptomics Problem Solving Environment with the correct tools. A workflow tool will be part of that.   In the two diagrams, the Genius environment for transcriptomics

studies, as developed within VL-e, is presented. Although it was used by the bioinformatics group within VL-e for prokaryote genomic studies, parts can also be used for eukaryotic transcriptomics research.

On the left hand side of both diagrams the general description of the type of software is illustrated. While on the right hand side the diagram shows some examples of tools that are present in that particular abstraction layer.

On the first diagram we can observe that, in order to enable these studies, a new abstraction layer is introduced being a domain generic e-Bioscience layer. This is a layer that's generic for a specific life-science e-science research domain. In this par-

ticular case, the necessary generic tools for prokaryotic transcriptomics studies are already present

http://www.microarray.nl/projects_inf.html.

But it is easy to imagine that when an imaging detector like a functional Magnetic Resonance Imaging (MRI) detector would be a requirement for a study, imaging tools should also become part of this domain generic layer.

In the second diagram (to the right) the VL-e functional model is presented again but now with the one extra e-Bioscience layer added to the stack of abstraction layers.

We have chosen workflow tools as an example for the diagrams. However, the layers also will contain other

kinds of tools like the imaging tools mentioned above. In both the illustrations Taverna is a workflow environment. "R" is a workflow tool that you plug into Taverna. CDF merger is a generic workflow, that can be used in many environments for transcriptomics. The CDF merger is also one of the tools that is used to get the correct and up-to-date annotation of analyzed genes and related pathways involved in biological processes.

http://www.systemsbiology.nl/datgen/transcriptomics/transcriptomics.html

Finally we have VLAM which is a workflow environment we developed before starting current VL-e project.

The illustration shows the advantage of the modularity we

have brought into our approach. If you would like to build something for proteomics or for metabolomics, you could use already a considerable number of the tools from the illustration and add domain specific and additional domain generic tools as they are needed.

For each (sub)domain in the life sciences, you want to have a community that will reuse tools whenever possible. These are present in the domain generic layer, called e-Bioscience in which the e-Bioscience tools reside. These tools can be used to develop Problem Solving Environments (PSEs) for different e-Bioscience subdomains Examples include Biobanking (the environment necessary to do cohort studies in medicine and biology), biodiversity (the study of the developments and presence of different species), e-Green Genetics (the study of plant breeding), and e-Food (the study of for instance why is something tasting bitter or the study of food safety) All these application subdomains are presented in the illustration. We believe that in future you will discover that there are many more generic functions that can be reused than you expect.

COOK Report: What you seem to be saying then is that the domain specific layer in the GENIUS slide above
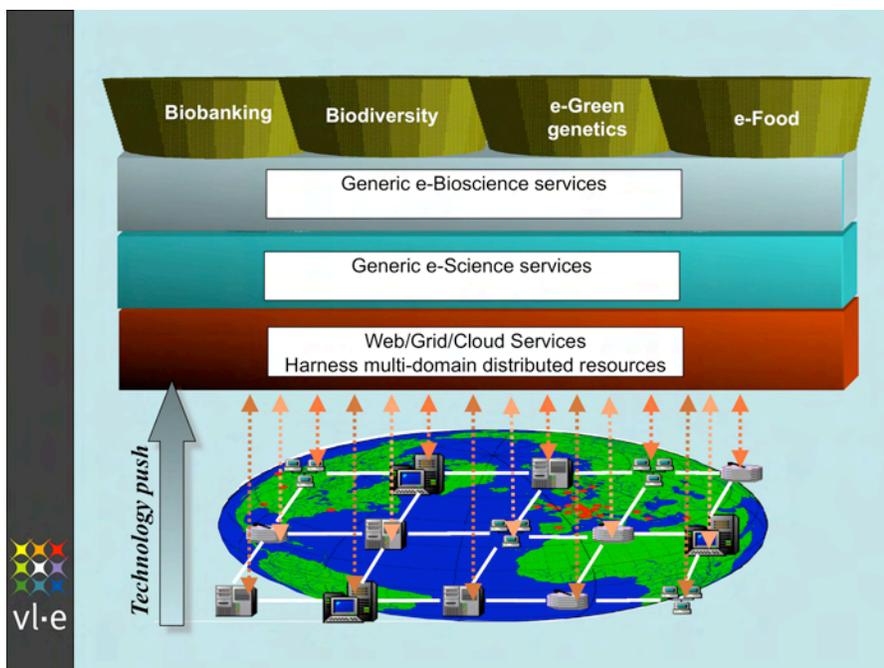
may be repeated in the diagram below to branch off further upstream into more detailed research tributaries using slightly different tools and plug ins into these other fours areas below?

Hertzberger: Yes, and what I am also trying to say is that within such an area you will find new subdomains that are also reasonably generic. The specific parts will be pushed higher and higher up, and the reuse of software will be pushed further and further along as well. But this is not something that I can prove at present.

COOK Report: But you *expect* to discover that as you find generic components in your new research they become reusable in more and more operations.

Hertzberger: Yes, and the reason why I think I am right is that the more one comes to understand about new e-science based experiments, the more we find additional generic aspects to those experiments than realized before. And of course, in such cases I also have the option of using software components already designed by others.

In addition functionality can be added to make them better suitable to the problem at hand. In other words, we can make them more generic. The

more a tool is applied for different research to solve slightly different problems, the more generic and powerful it becomes. This is why I view myself as being halfway between being a computer scientist and an experimental scientist. A computer scientist wants to build as many reusable tools as possible. An experimentalist says I have to solve my problem even if it means risking reinventing the wheel. And by the way, when you make the tools modular, they become easier to re-use, because they can be applied to solve different problems.

The more modular you make your software, the more successful you will be.

There is also a philosophical point here that seems essential. If you are going to use a computer in your experimentation, (and this will happen more and more,) you must ask whether the methodology (design) of the experimentation is not determined by the fact that you use that computer. This implies that e-science is just another form of science and not only an enabler making domain science possible.

Apart from this, one can also learn from other fields of

automation (like office or administrative automation) that automating a domain process always requires a better understanding and consequently a rationalization of that process. This is a necessary first step towards further standardization of the domain process in order to make it suitable for automation.

Participating in e-science, research will force domain communities into better thinking through their research problems and consequently better understanding their core business. This is a first step towards rationalization and will help in further standardization of their research methodologies. This on its turn is a prerequisite for exchanging data/information an essential condition for collaborative research.

User Experience with e-Science

The biologist responsible for the development of the technology chain for microarray transcriptomics likes to say we are now going into the fourth phase of e-science development in the Netherlands. The first phase was when most application domains were just learning all the stuff we've mentioned above, and realizing slowly that they had to change their way of doing science. Regrettably, it has to be said that for the most, only

those scientists that were actively involved in the e-science developments currently understand the challenges we are facing. For most life-science scientists, the computer is still just an interface to rather simple software tools.

In the second phase, we learned about the layered structure and the fact that we needed to build Problem Solving Environments (PSEs) with generic elements. In this phase the e-science approach became immediately extremely complicated and frustrating.

In the third phase, we better understood that there was something as a domain-generic layer. And more importantly, we all realized our roles in the whole VL-e complex. We should work from the perspective of our own expertise, and our objectives and results should reflect that. For

applications this means that we have to produce functional PSEs in order to prove the e-science concept.

### *…experimentation is the only way to verify our ideas…*

Now we are indeed entering the fourth phase. Here, we have to not only build these PSEs, but also to support them in a structured way in order to make them successful. These PSEs are so complex that they need expert support. Otherwise life-science domain scientists actually doing the experiments cannot use them. So for all involved application scientists, this e-science quest has been a difficult experience, with little immediate reward in their application domain. However, no one even considers abandoning this approach, as the long-term rewards with

respect to scientific competitive advantages are so apparent. In this case inclusion of e-science in the basic curricula of biology and medicine will also help.

The VL-e Experimentation and Software Production Infrastructure

Hertzberger: As I mentioned earlier, for me the proof of the pudding is in the eating. Therefore we constructed two experimentation environments that acted as the software production and test infrastructure of VL-e. On the right hand side of the slide on the next page is the rapid prototyping environment for software development, while you will see on the left hand side the proof of concept environment designated to test out and get experience with application software. In addition the diagram shows a test and certification environment that helps in moving software from right hand experimental environment to the left hand proof of concept environment.

COOK Report: Therefore the VL-E Infrastructure diagram that follows is very important. It describes the process that you use to mature your software from an experimental environment into a certification one, and
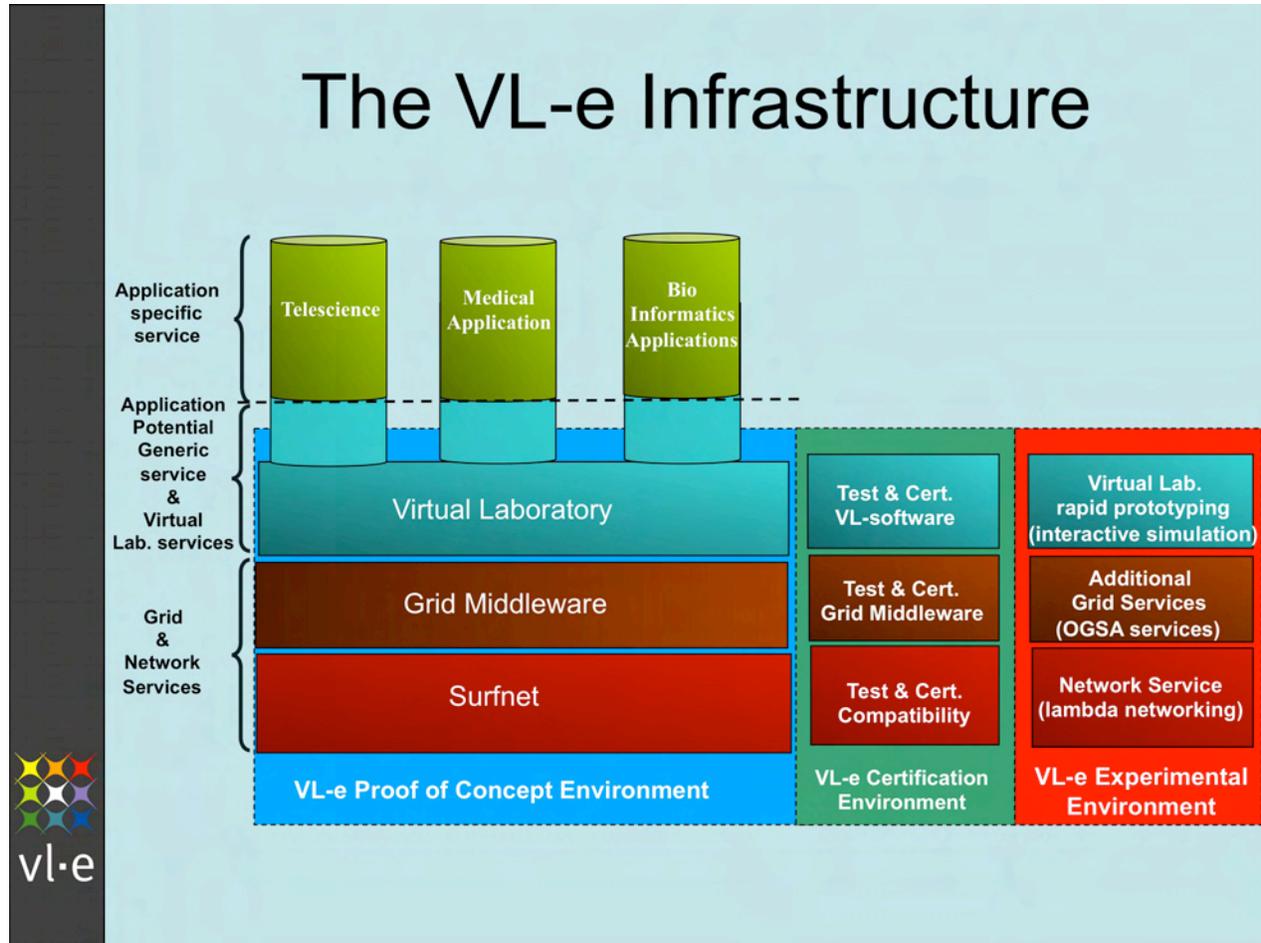
then into what you call proof of concept, before being released to first the community of different user in the VL-e project in the form of the proof of concept. Moreover the "e-Science Roll Out" slide 2 pages down shows how the software is further made available to user communities spread all over the Netherlands via the Grid based ICT research Netherlands infrastructure called BiG Grid.

Hertzberger: Yes. And the number of BiG Grid users is growing steadily.

COOK Report: Take me then on a tour through the draw-ings to detail how your process works.

Hertzberger: On the "VL-e Infrastructure" slide at the extreme right in the red box you see the words "VL-e experimental environment." This is our rapid prototyping environment in which scientists can develop their tools without having to take into account that the rest of VL-e exists. They can play.

COOK Report: This is what Google and other folk would call a "Sandbox" on the sense of the contained box of sand in which children play and build things and then tear them down.

Hertzberger: I see. We call it rapid prototyping, and this rapid prototyping is exclusively used by computer scientists. The current system (DAS-3) is a distributed system with multiple 10 Gb/s links (lambdas) between the clusters, provided by SURFnet.

It provides an excellent experimentation environment for computer scientists. Together with the application people, they can decide on a tool that they want to enhance, or they can build a new tool.
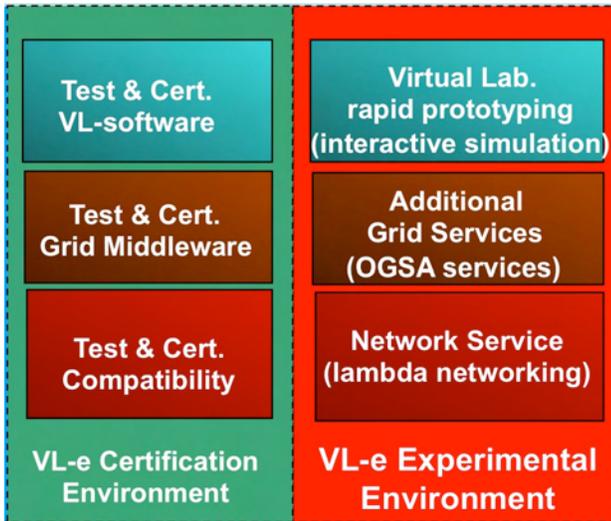
Let's assume that the computer scientists and application scientists want to develop a tool to use in medical applications. They take the tools available to them in the right hand red column, and they use them to develop their new tool.

Then at a certain moment they are happy. They say, "OK this tool can now be used in our medical application user environment."

COOK Report: In other words they release it to medical scientist end users to play with?

Hertzberger: Yes, but be-

fore they can do this they first must go through the test and certification environment that is shown in green above on the left side. The reason is that the tool now has to become an operational part of a larger software environment.

COOK Report: The certification process is to see how it fits with all the other pieces of the environment.

Hertzberger: Correct. It must fit into and interoperate with the virtual laboratory tools and the grid middleware. When it passes certification it comes into proof of concept on the left hand side, which completes the migration process. The scientists are using it, and based on their experience, we decide whether it needs modification.
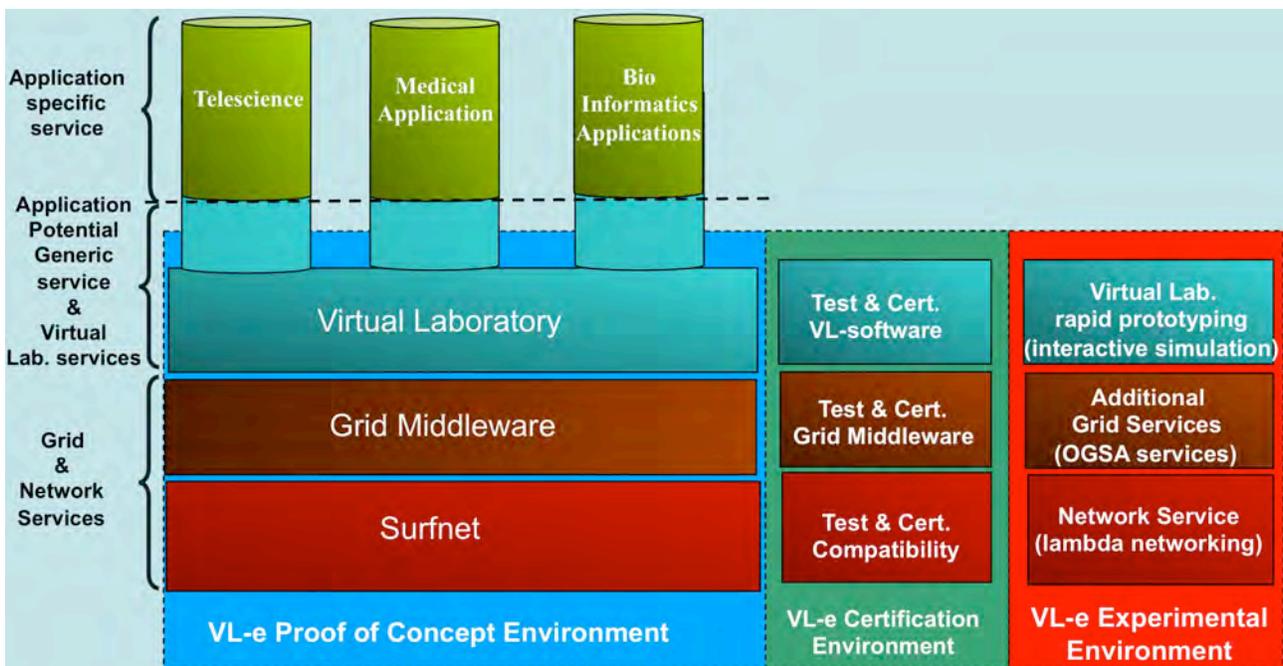
By the way, the three green and turquoise cylinders that you see sitting on the Virtual Laboratory on the left side, conceptually should also be shown sitting on the red column at the right as well as on the green column in the center. But that would expand the size of the diagram too much.
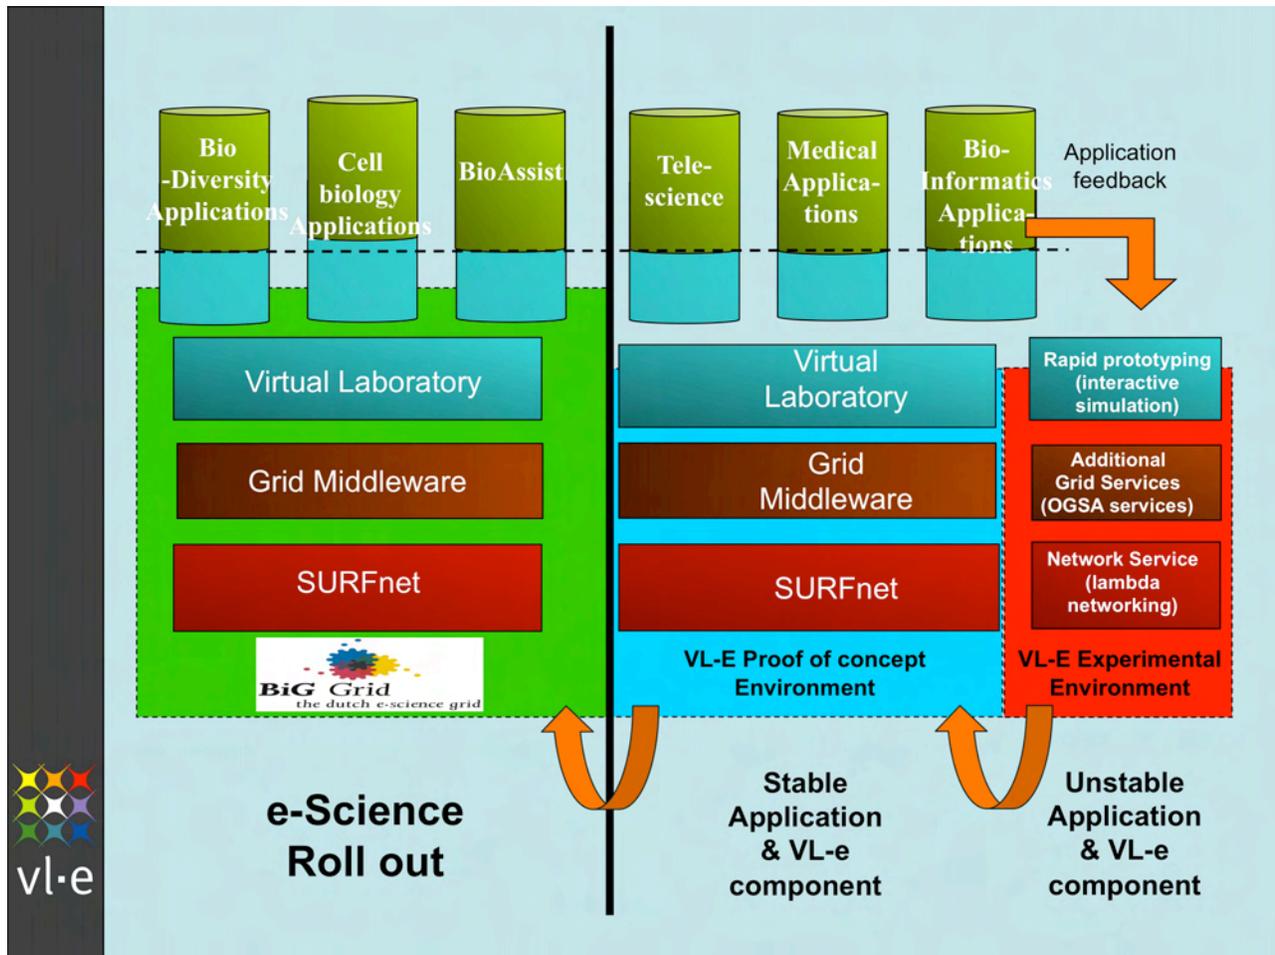
We developed something called the V-browser that went through this process a number of times before the users found it acceptable. This V-browser is a tool designed to make the study of medical information easier and is now being used in medical communities. Another example is the Java-GAT, which hides the different underlying grid middleware. JavaGAT is also used by other international projects (for example TeraGrid and D-Grid) and has also led to the SAGA standardization effort.

For VL-e the proof of concept is part of the proof of the pudding and the proof of the pudding is – as I have said – in the eating. If nothing else this should have been what we promised to deliver to our government. That there was some difficulty in understanding what we did is illustrated

…transferring our proof-of-concept software went well, most of it ran without modification…

by the fact that people from government agencies monitoring our project in the beginning complained that this activity was reducing the amount of manpower that could be devoted to research. Now let's look at the words "e-science roll-out" in the left hand column of the slide that shows the entire process.

COOK Report: The eating of the pudding is found in the turquoise colored proof of concept above?

Hertzberger: Yes. And when the eating of the pudding was declared good by the VL-e user community, it

was transferred to BiG Grid where it is used by about 40 user communities now. BiG Grid is a separate project organization based on a proposal that was submitted by Nikhef (the Dutch institute for subatomic physics), the Netherlands Bio-Informatics Center (NBIC) and the National Computing Facilities foundation (NCF). In 2006 the project was awarded 29 million euro by NWO. BiG Grid was established in order to set up a grid based user infrastructure for our scientists. The Dutch National Computer Center SARA was brought on board as an operational partner providing grid operations and support to the BiG Grid communities. One major part of the BiG Grid activities is the running of the Dutch Tier I facility for LHC, which is operated jointly by Nikhef and SARA.

Generic software should be understood to be an essential part of basic infrastructure

COOK Report: In other words BiG Grid is the railroad tracks on which VL-e's e-science middleware runs?

Hertzberger: Exactly. You have seen that BiG Grid is a

separate organization with its own funding that will run for another two to three years. It will support other applications besides those of VL-e. And those folks will make new applications that run on top of ours.   In transferring our proof of concept software to BiG Grid, it turned out that most of it ran without the need for any modification.

COOK Report: So what you are saying is that the middle bars, Grid middleware, and virtual laboratory transferred from the VL-e proof of concept into the e-science roll out within BiG Grid. So you handed things over from VL-e to BiG Grid as direct work on VL-e stopped?

Hertzberger: Yes, the proof of concept software will be further developed, but at this point it will be done within BiG Grid.

COOK Report: And in 2010 the e-science Research Center that replaces VL-e will be free to do whatever people

want it to do?

Hertzberger: Correct. And they may even use a different model. Let me conclude with the description of the VL-e experimentation environments. The proof of concept now resides with BiG Grid while the rapid prototyping environment still runs on the Distributed ASCI Supercomputer (DAS) of the computer science research school ASCI. Big Grid can use the Test and Certification environment in the future for its own software.

…funding is not unlimited… decision-makers had better invest in infrastructure first.

This kind of development mechanism was always an idea I had and strived for. I did not expect the migration
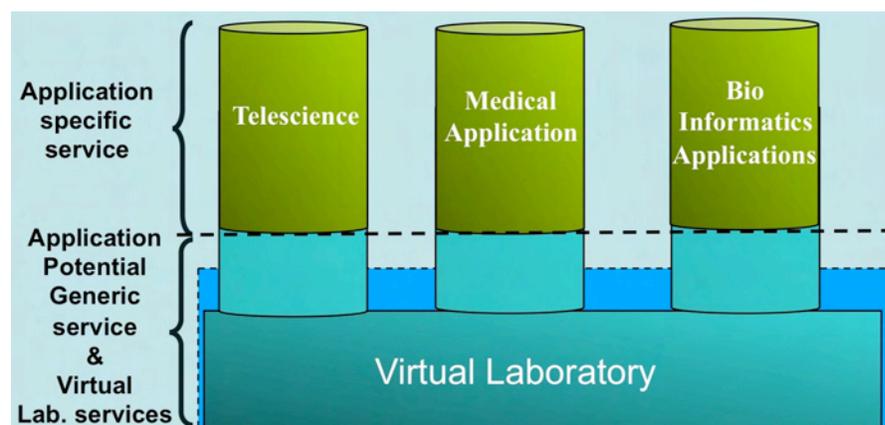
to run so smoothly, and those responsible for it did a wonderful job. But on the other hand, we did have a problem because a lot of the tools in some of the tool boxes are not part of the proof of concept environment. The domain generic tools are represented by turquoise color bottoms of the cylinders below. But note that as the color shows, they overlap and extend into the virtual laboratory as well.

The reason is that the validation process was designed to make software supportable in an operational ICT environment. Whereas that approach is essential for the base infrastructure software, it is less well suited to software and toolboxes supporting the core scientific process. Here, scientists were always making changes to their tool boxes, because they got new ideas. When they do they must adapt their tool boxes.

What we are doing now is to cope with this problem by, on the one hand, making tools as generic as we can. But on the other hand, since many are 'domain generic', a more agile deployment of toolboxes based on community-supported software distribution has proven more effective. In BiG Grid a team of e-science support and development experts helps researchers to achieve this agile environment built on top

of the base software suites. As a scientist you need to have tool flexibility. And the problem was that changes to the tool boxes also impacted the ability of the toolbox to interface with the virtual laboratory services.

The key question for us was how do we support the tools? The answer was that support came primarily from people in BiG Grid. What I am trying to explain is that the structure we have built is more flexible than it looks to be in these figures, which give you just the impression of what is generic and what is application specific. In actual practice what happens is that someone builds his own tools and then asks for support when those are used by a community.

COOK Report: It sounds like the boundary between virtual laboratory and domain generic tools is fuzzy?

Hertzberger: It is indeed fuzzy. This fuzziness first seemed to be a failure, but is in fact a big success. Sometimes things are running in the virtual laboratory without having obtained certification from the BiG Grid organization supporting it as an operational service. This is possible because the user community gives the necessary support. Therefore the users are happy. Then you have, as a result, a lot more flexibility.

Because if a user discovers something sufficiently generic, the engineer supporting him can integrate it into the toolbox he is using. These are things that we did not try to design in advance but learned as the project went forward. In practice we have either generic tools and especially domain generic tools, or we have application specific tools. And these run in the BiG Grid environment. Depending on which community you are talking about they will use certain toolboxes and certain tools. The more modular you make your software the more successful you will be.

The Economic Impact

The Committee of Wise Men in the Netherlands was responsible for giving us 20 million euros and we had to find another plus 20 million euros matching. We had to define targets. The most important was to show that we had impact at the end.

What is illustrated with this last slide is that in some of the cases we built applications and tools, which if you had to build them from scratch, would already have cost you more than our total project budget. I will give you an example. In VL-e we built a virtual lab for ecological analysis based on our generic

tools and some specific ones. This is now in use by municipalities, project developers and building companies in the Netherlands. Before they can start construction activities, they have to make an inventory of rare plants and animals living on the land where they wish to build. They can use shared data and methodologies or tools to find out what species there are and what they are allowed to do or not to do with each. Currently about 50 people are working on this project with a total budget of 20 million

Euros. This is presented as the spill over explicitly called out in the Economic impact slide because the roots are all in VL-e.

Another tool that we built and is now in actual operation is a bird avoidance system for our air force. By combining a sensor network and models, it predicts where flocks of birds are going and advises pilots how to avoid them. The same system is also used for scientific research into migration patterns of birds. Developing these tools in VL-e had

the advantage that it was not necessary to start from scratch. We had domain as well as e-science knowledge available. Moreover, we possessed the necessary e-science (proof of concept) environment on which to build.

COOK Report: Are the figures on the right side – 20, 10 and 15 adding up to 45 the projected costs of doing those tool sets from scratch?

Hertzberger: Yes and some people claim that I was far too conservative and that the actual from start costs would be two or three times higher than my figures. Somebody retired from IBM estimated the cost would be more than three times my figures. But I am a scientist so I chose conservatively.

COOK Report: Because you have leverage – a multiplier effect?

Hertzberger: I'm happy to say, yes, we have a multiplier effect.

VI. Growing E-Science Domains for The Netherlands

Roadmap for a Next Generation of Science

"Dutch e-Science efforts are at the leading edge"

While the development of e-science and related infrastructure is a global effort, Dutch scientists have been at the leading edge of these developments. The VL-e mid-term review Committee concluded in its report:

"The Committee considers VL-e to be an excellent project at its mid-life point with high quality contributions. It has pushed back the frontiers of e-science, whose aim is to induce a paradigm shift in the methodology of science itself, a difficult and worthy challenge."

As demonstrated by the cases in this section, the infrastructure elements needed by a wide variety of sciences are very similar. All science cases require a rich landscape of resources, consisting of networks, computing, visualization and storage, but foremost services that are instrumental in integrating this infrastructure and that enable the scientist to effectively address and use it for his research. The most efficient solution to this e-science challenge is to iden-

tify those common features and deploy one infrastructure that takes care of these common problems

[**Editor's note**: the e-science interview Chapter V above makes this point. Note also that the following abbreviated catalogue of Dutch e-science effort is a shortened summary of an extensive department store like inventory. Here are top level highlights. A text box follows with subdivided areas in life sciences and physical sciences.]

Life sciences

Biodiversity
A variety of organizations collect samples of species and make observations of their occurrence and abundance. These collections serve not only the study of agriculture, water quality, bio- prospecting, conservation, but also research in systematics, biogeography and population ecology. These data are more and more appearing electronically, scattered around various labs and institutes, whilst also massive amounts of spatially explicit data on the collection sites have become available. [. . . .]

Editor's note: – Biodiversity was just the first topic listed here. The remaining are Biobanking, Metabolomics Biomolecular informatics, Cognitive science and have been moved to the text box below.

Physics, astronomy, chemistry

### Molecular simulation
Molecular Simulation has the potential to play a central role in the design of new materials and processes, and in the modeling of biological processes. [snip] Editor's note – Molecular simulation was just the first topic listed here. The remaining are Particle physics, Astronomy Catalysis for energy conservation, Fluid dynamics and have been moved to the text box below.

Climate research

Climate induced weather extremes
The consequences of global warming for Western Europe include an increased probability for hot summers and higher winter precipitation, the latter leading to an increased amount of water to be transported through Dutch rivers to the North Sea, challenging the Netherlands for
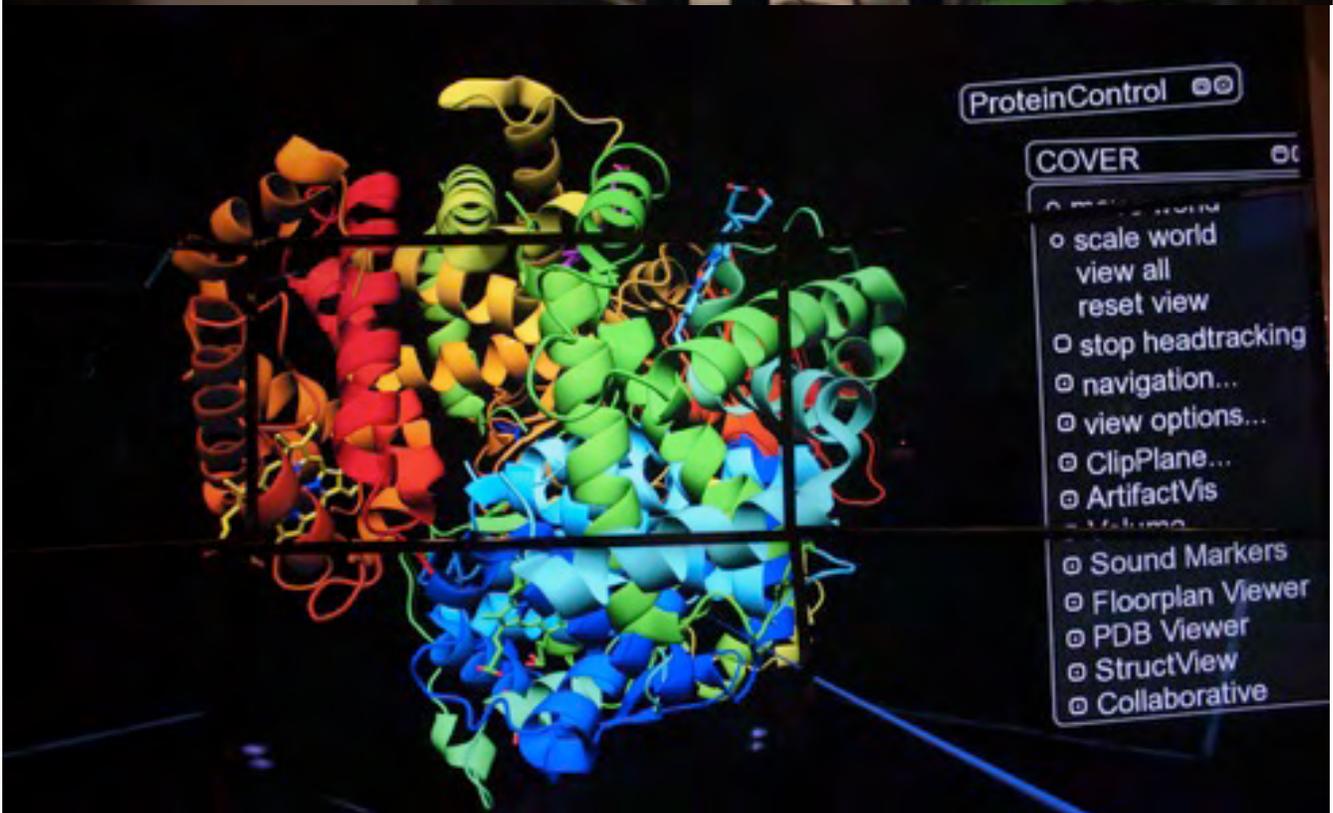
These photos from The KAUST Exhibit at Supercomputing 09 indicate the use of 3d visualization tools where two scientists can explore protein folding. In this instance they walk toward the display with a visor and a pointer that enables them to steer the display. They can communicate in real time examining the same model from different perspectives. The second picture shows a virtual navigation menu by means of which the user on the show floor can change his interaction with the display which uses eleven "tiled" LCD monitors, nine of which are visible.

decades to come. It requires a combined investigation of the underlying atmospheric and climate processes. Simulations with an atmosphere model with a horizontal resolution of 400 km are possible, however, it is necessary to go to at least 40 km which will require a more than 100-fold (10*10) increase in required computing time. The same factor 10 holds with respect to ocean model resolution. For better accuracy, we even need to go beyond that.

Arts and Humanities

Services for the e-humanities

Life Science Detail

Biobanking

The development of a (federated) biobanking database infrastructure as a critical component for biomedical life sciences research and discovery is urgently needed. The explosion of genotypic and phenotypic data requires that data are properly stored, accessed, managed, queried, analyzed, and shared with others. Flexible data integration is needed, both across locations, and across application platforms, varying from genetic, clinical, molecular, and demographical data, to analytical and biostatistics data [snip]

Metabolomics

Metabolite profiles reflect biochemical changes, which should be properly identified and analyzed such that the different type of variations, i.e. the changes due to disease, nutrition, life style, etc, can be differentiated. However, the field of metabolomics is challenged with the huge amounts of high-resolution mass spectrometric data produced and the lack of an advanced bioinformatics framework for proper handling and pre-processing these data. [snip]

Biomolecular informatics

The Center for Molecular and Biomolecular Informatics (CMBI) has long been involved in projects aimed at improving the precision and accuracy of three-dimensional coordinates of macromolecular structures. This is an important aspect of a larger project aimed at the elucidation of the molecular origin of genetic disorders. Late 2007 a major step was set when all protein structure files that were ever solved worldwide were re-refined using modern software and today's understanding of geometric and energetic aspects of protein structures. This was a major operation that involved hundreds of computers spread over Europe but combined in the EMBRACE Virtual Organization for biosciences.

Cognitive science

It is widely believed that the structure of the human visual systems is related to (statistical) regularities of the visual world. Understanding and simulating the human visual systems therefore requires understanding these regularities. . To this end, cognitive scientists from the University of Amsterdam study properties of the visual world and their influence on the structure of the human visual system using whole-brain modalities such as M/EEG and fMRI [snip]

Physics, Astronomy, Chemistry Detail

Particle physics

Particle physics has made striking advances in describing the intimate structure of matter and the forces that determine the architecture of the universe. Nevertheless, fundamental questions like: 'What is the origin of mass?' and 'What happened to anti-matter since the Big Bang?' remain. Parts of the answer to these and other questions are likely to come from the Large Hadron Collider (LHC) at CERN. [snip]

Astronomy

LOFAR is an innovative effort to force a breakthrough in sensitivity for astronomical observations at low radio frequencies. It will explore the low frequency (10-200 MHz) radio sky at high resolution for the very first time, focusing on several scientific goals: the study of the very early universe, the study of exotic phenomena such as the formation of massive black holes, and (clusters of) galaxies, observation of Gamma Ray Bursts, radio supernovae, intermediate black holes, flare stars and exo-planets, and the detection of extremely high-energy particles.

Catalysis for energy conservation

On a global scale, energy conservation and ecology-friendly industrial production processes have captured much attention. Examples are the development of effective detergents at low temperatures and the development of low-energy processes in the oil industry. An important area in chemical science copes with these phenomena and problems. [snip]

Fluid dynamics

Fluid dynamics is the scientific discipline that covers all kinds of flow phenomena: hydrodynamics, aerodynamics, climate research, meteorology, combustion processes, etc. With the presence of turbulence, flow processes in general are extremely complex. [snip]

Digital humanities will profit from infrastructural support for the handling of multimedia content in the interpretative processes that are the heart of humanities research. In order to allow individual scholars and groups to share source data and the annotations that are the result of interpretation or comparison, the humanities need advanced collaboratories: interoperable platforms that can support and integrate content analysis (e.g. (semi)-automated transcription of spoken word collections, semantic clustering, semantic video annotation), that stimulates community building and that can be coupled to visualization tools.[snip]

Large digital collections of annotated speech and other language data (text, audio and video), data bases on population characteristics (census data, population reg- isters), economic time series, and geospatial data require new approaches in fields such as phonetics, sociology, econometrics, demography, history, geography and archaeology. DANS, the national center for permanent access to research data for the humanities and social sciences, is exploring the potential of the data grid and the linkage of heterogeneous data resources."

VII. Potential Customers with Global Agendas

An interview with David Zakim, MD

Dr. David Zakim wants globally available multi-disciplinary databases to help improve individual patient care in real-time.

David Zakim is the kind of "customer" that will increasingly need the tools that Bob Hertzberger's team is developing in the VL-e and the research infrastructure described earlier in this report. Zakim has a big agenda. He wants to transform the practice of clinical medicine, which cannot be accomplished without leveraging e-science.

In 2000 David Zakim retired from Cornell Medical School after a 40 year career in academic medicine. But he was unable to stop thinking about the chronic problems that limited the ability of physicians to deliver effective care to their patients.

Zakim knew that physicians rely on memory to treat patients. The physician pretends he or she can carry around in memory all the medical knowledge needed to match presenting problems and formulate a defensible diagnosis and treatment plan. This is a physical impossibility because of the enormous scope and complexity of the knowledge base for practice.

But whereas there is already too much to learn, there is too often insufficient knowledge for accurate prediction of risks and outcomes. Thus, everyday medical practice struggles with two significant knowledge problems: too much to learn but not enough to practice efficiently.

What's missing, Zakim realized, is technology that insures the application of relevant evidence to management of clinical problems on a day-to-day basis and a clinical research infrastructure to support better understanding of risk factors for disease and the therapeutic efficacy of a variety of treatments.

These are not new problems. 30 years ago I published an article in the June 1979 Futurist about the early work of Dr. Lawrence Weed on applying computers to these medical issues. Not much has changed in practice of medicine, however, including that it consumes an ever-increasing percentage of GDP. But technology has leapt ahead.

Zakim started work on these issues upon retirement from Cornell and founded the Institute for Digital Medicine in 2007 as the organization to fulfill his vision for improving the quality of health care and reinvigorating clinical research.

http://www.idm-foundation.org

Medicine is practiced on the false premise that the physician can hold all the knowledge he will ever need in memory

David Zakim MD

Zakim's CLEOS system is designed to change dramatically the acquisition of patient histories, the assembly of patient medical records, the coordination of all personal medical data gathered over a

patient's lifetime and the assembly of very large clinical databases that describe the behavior of disease across time.

...Zakim realized what was missing: real-time databases that instantly and continually improve patient management

Zakim's goals are three-fold: enable a computer system to collect, analyze and make available the broadest possible range of patient data, from patients around the world via direct interactions with people seeking health care; assemble large clinical databases that are amenable to mathematical analysis; and allow delivery of management advice through real-time comparisons of data for individual patients with the accumulated clinical experience within a world-wide clinical database.

In cooperation with Dr. Mark Dominik Alscher, Professor of Internal Medicine, the University of Tuebingen, Germany and director, department of Internal Medicine and Nephrology, Robert Bosch Krankenhaus (RBK), Germany, Dr. Zakim has

been implementing a computerized questionnaire and decision-tree-based system for collecting and interpreting clinical data via direct interaction with patients in the absence of physician input. That work is described in the following article in the literature.

http://www.biomedcentral.com/1472-6947/8/50

The real-time capabilities for clinicians envisioned by David Zakim, however, will require a global ICT research infrastructure. It will need software tools for multi-disciplinary teams, capacity for massive databases, fed by millions of heterogeneous sources, well-oiled database integration, with robust and agile networks. All this will necessitate massive real-time processing and data transport capabilities.

Zakim needs specifically the kind of e-science tools, teams and work processes that Kees Neggers, Bob Hertzberger and their many colleagues have spent years evolving.

Recently Zakim was emailed an excerpt from the presentation Kees Neggers gave at NORDUnet in September, 2009, plus a brief overview of the goals, strategies and achievements of ICT research in the Netherlands.

"Who is Neggers?" Zakim emailed back. "And how fast can I get to meet him? His paper describes the problems and aims of exactly what I

am trying to achieve in a specific knowledge domain."

[**Editor's note:** At the start of his presentation Neggers describes the SURF Roadmap Towards a National ICT Research Infrastructure developed jointly by the ICT Infrastructure providers in The Netherlands described earlier in this report].

Most of this book has focused on the practitioners, processes and "products" of advanced ICT infrastructure research in the Netherlands. Before going to press, it seemed a good idea to spend some time talking to a potential user as well.

Here are highlights of a January 27 follow up to an earlier two hour interview with Dr. David Zakim.

COOK Report: David, what is your reaction to Kees Neggers presentation on what SURFnet and Bob Hertzberger are doing.

Who is Neggers and how fast can I meet him?

Zakim: The first of Kees Neggers' recent research trends is system level science, which is defined as the integration of diverse sources of knowledge about the constituent parts of a complex system. This describes some

areas of science, but it also applies to the entirety of medicine. Our problem is a pediatrician knows very little about adult medicine; internists essentially know zero pediatrics. An obstetrician can't take care of a patient with heart disease, nor can a cardiologist take care of a pregnant woman. Even within the specialties of adult medicine, such as internal medicine, a gastroenterologist will know little about the diagnosis and management of cardiac disease.

COOK Report: The human being is a complex system.

Zakim: Enormously complex. Because of that complexity, and because of the *scope* of knowledge available, medicine has been increasingly cut into smaller and smaller specialties and subspecialties. Specialists do not communicate with each other; they just refer patients.

Medicine, in Neggers' terms, is system level science. To capture the knowledge base, to enable a machine to help an individual to practice medicine, one needs to integrate knowledge from diverse groups of physicians with different subspecialties, different knowledge sets, and a wide range clinical experience.

The second point Neggers

made is the importance of multi-disciplinary research. Physicians are already in some form doing multi-disciplinary research already. Epidemiologists, who don't treat patients, use MDs examine large databases to discern cause and effect.

We need a far more robust form of multi-disciplinary research. In clinical research we need non-medical expertise such as computational expertise and mathematical expertise. Within the discipline of medicine, examination of clinical data is based exclusively on concepts of patho-physiology. (Neggers refers to exploring data sets with *upfront* hypotheses.) But we know the power of that has reached its limits.

COOK Report: Doesn't an upfront hypotheses bias your thinking and choices?

…powerful algorithms can analyze databases of unlimited size.

Zakim: More than bias, it says we're going to look for a relationship between this set of data and some outcome. The dataset is cherry-picked according to pre-conceived notions of what kinds of events might impact the out-

come of interest.

We *know* smoking contributes to coronary disease. So we might decide that algorithms for predicting the incidence of coronary events need to be improved. Consequently we're going to sub-define smoking as one way of determining a better algorithm. Now that excludes all other data and is based on the idea of what we know, i.e. we *know* smoking is important. But it could be that eye color is important as well. No one has ever checked.

It's impossible to look at all possibilities, except using patho-physiology, unless you posit the question that we have to begin to look with abstract mathematical algorithms. In that case we need to involve mathematicians in examining very large databases. They can examine databases of unlimited size and do it with powerful algorithms that don't depend on pre-formed ideas of cause and effect. Almost all epidemiology is based on the idea that X causes Y, and then prove it or not prove it.

Let's say we have a factor with 10 possibilities. And none of them appear to be important. None of the combinations turn out to be significant. But that doesn't mean that one of them is not

important. It's simply that you couldn't show their importance in the context of how you examined the data. They might, however, turn out to be important in the context of 100 or 1,000 data fields when relating to the outcome, let's say, of a coronary event.

The CLEOS system developed at our Institute for Digital medicine allows us to populate "yes or no" in approximately 18,000 to 19,000 data fields. Mathematicians are skilled in examining that kind of data situation. Therefore we have to bring to bear the power of mathematics in developing predictive algorithms for medicine.

Multi-disciplinary research cannot be conducted the way clinical research and epidemiology *have* been conducted. A multi-disciplinary approach depends on building very large and standardized clinical databases that accrue data on individual patients across time. But we also need the ability to integrate data from *millions* of patients across time. This requires an extremely powerful computational infrastructure. You need very smart people to build and maintain the database structure, to build a networr, and to integrate all of this information. I know that this is true. There is nothing available now that can sup-

port that sort of approach.

COOK Report: Why does the Netherlands' system appear so promising?

Zakim: In Netherland they have *set out* to develop software that *can* support that sort of approach. No one else that I know has that. I read about a company in North Carolina that builds most of inventory databases for big corporations. They're innovative but of course their approach is proprietary.

COOK Report: You've made it clear in your papers that the proprietary nature of the software would preclude it. Under "Representing the public interest in health care", you wrote **"***CLEOS®* is not a commercial product". It is owned by the IDM Foundation because the software can be corrupted easily for narrow commercial gain, as for example recommending products of a single company or always recommending patent-protected medications versus generics.

"An equally important reason for non-commercial ownership of *CLEOS* lies in the generation of otherwise unavailable clinical databases. This power, in commercial hands, would privatize medical knowledge to make new medical knowledge a profit center. And commercial

ownership of programs like *CLEOS* would transfer development and deployment of new medical knowledge from the medical profession to profit-driven entrepreneurs, who would be guided by 'share-holder value' not the public interest."

Zakim: The third thing Neggers emphasizes is real-time processing of very large datasets. At the current time there is no real-time processing of medical databases. Typically if a databases exists, someone will spend a year testing the information in the database according to some hypothesis. He or she write a paper and submit it to the *Annals of Internal Medicine*. Two or three years after that the paper gets published. Assuming it gets read and assuming it's right, it's still up to an individual physician to integrate the material in that paper into his or her practice.

COOK Report: What do you mean by "real-time processing"?

Zakim: What I mean by real-time processing, and what I interpreted Neggers's meaning to be, and why it hit me, is this. Miss Smith comes to your office. She suffers from diabetes and is age 30. With real-time processing, whatever entry Miss Smith adds to her fundamental database, in her digitized medical records

within CLEOS, is compared with *every other patient in the database in real time*. The output decision as to how Miss Smith's management should be changed or left alone - because of how Miss Smith's parameters compare to someone who had diabetes at 30 and... then had X event at 40 or at 60.

It gets more and more predictive.

So Miss Smith is compared with... let's say we have 50 million people in the database. In real time, Miss Smith's clinical parameters are compared with 50 million other people in the database to determine what is the best way to manage Miss Smith on January 26, 2010, at age 30 with the following diagnoses and co-morbid state, if she takes insulin, and this and that, and something from hypertension, and has a creatine of 3.6. The system recommends a management decision by the time Miss Smith finishes entering her data. *That* is real time processing.

Physicians must be able to directly input the knowledge themselves, or it will not work.

COOK Report: That's awesome.

Zakim: That's how it's *supposed* to work. We call that the doctor making decisions on the basis of clinical experience. But the literature suggests the doctor's clinical experience is defined by the last patient he or she remembers who had a similar illness.

Let's say we have 50 million people in the database. Almost 8 million will have diabetes. Maybe 600,000 had diabetes at 30. So the database is cut to 600,000. All of that is done instantly. The system matches Miss Smith with all present and past patients whose data can help the clinician recommend the best course of management for Miss Smith's diabetes, at that moment, in real time.

Do you know how many participants were in the Framingham Heart Study, which is the largest predictive database in our universe right now? It started with 8,000 and later they got 12,000 more. And it's an enormously powerful database. But it does not have all of the data fields in it. And it doesn't follow the patients as carefully as an electronic system like CLEOS can. So imagine what you could do with a database of 600,000.

COOK Report: Over time as you accumulate patients and data, the system should get more and more powerful.

Zakim: Absolutely. It gets more and more predictive.

COOK Report: How you intend to start this up over the next 2 to 5 years within the context of an infrastructure such as the Netherlands'?

Zakim: First, I don't know the details of what it is. Second, I'd have to know the details, not the nitty gritty, in a conceptual way. We'd need to have a discussion between those people like myself who know about medicine, who know what's in CLEOS, who know what issues are involved in adding knowledge to CLEOS. And I would want a mathematician who has been very involved in, for example, analysis of micro-array data, which is closely analogous to the problem of analyzing the data CLEOS collects.

COOK Report: I have the impression you're starting with a few hundred people, then a few thousand people, and now you are doing some data accumulation in Germany. And at this point it's focused on cardio-pulmonary data. Are you actually writing code that takes the patient questionnaire and follows it down the diagnostic

path where you begin to pull in specific medical literature?

Zakim: I don't write code. I have a graphic user interface. I write the code as I add knowledge. Now this is a very important point, because you will *not* successfully build this sort of a system if you put knowledge engineers between the system and the physicians. You will go through hundreds of millions of dollars and you'll come up with garbage. Physicians have got to be able to work at an inter-face so that they directly input the knowledge themselves. It will *not* work otherwise. I *know* that.

COOK Report: What would be an example?

Zakim: When you come to my office and I ask you only *one* question that is not responsive to something you're telling me—which is: "What's wrong? Why have you come today?" Now everything I ask subsequent to that is reactive. I interpret your answer in the context of everything you've told me, and then I ask you the next question.

When I write this out, I'm working in an abstract situation. I have no responses; I have to put in all possibilities. So I will write it and I will incorporate it. Then when I play it back, which is looking at the interview, I can say,

Oh, boy. How stupid could I have been? I left something out. The patient won't notice the sequence but it makes it very difficult to ask follow-up questions. So it's a *re-iterative process* in which you enter knowledge, and you have to play the knowledge back to yourself to refine the knowledge.

…not just very large data sets, but enormous data sets…

Okay, now I work with a knowledge engineer, and I can tell the knowledge engineer I want him to program this set of questions in this sequence according to this set of answers. And then the knowledge engineer sends me an email saying, Okay, you can see it.

COOK Report: I'm guessing from what you're saying, that there's much tree branching and perhaps fractals?

Zakim: There are some fractals in that I reuse nests of questions. For example, fever. Completely unspecific. It comes up over and over again. It will be used in multiple places in a medical history, in a million medical histories, in more than a million different places.

COOK Report: Is it tree branching that's not predictable?

Zakim: No, the trees are completely determined. However, there are more pathways through the trees than there are people on this planet.

COOK Report: And the knowledge engineer is not aware of the issue of those pathways?

Zakim: Correct. He doesn't know where to put what. Not only that, it will take too long to feed it back to him. And the ability to change things on the fly and then see what that did to it...that's a very powerful experience. I've learned that I can build logic structures that I never see when I'm building it. I only see it later.

COOK Report: Is this one of the example of the mathematicians and mathematical algorithms?

Zakim: No. The application of mathematics in Neggers' view is the issue of multi-disciplinary research, but also the combination of mathematical analysis with the computational infrastructure in order to provide real-time processing... not of just very large data sets, but of *enormous* data sets... worldwide data sets.

COOK Report: So you have already developed a pretty good questionnaire and patient history form in your hospital in Stuttgart. But when you want to take accumulations of this patient data, and when you want to complete more data in order to do new kinds of clinical research that has been heretofore impossible... that's when you begin to need the kind of system we've been talking about in the Netherlands?

Zakim: That is correct. But it goes further than that. Because every time data are added to the database... every time Miss Smith adds data to the database... her data can potentially impact *every other* diabetic in the database. The knowledge of what happened to her in response to this or that (or what didn't happen) could impact everyone else in the database.

...it's very hard to build infrastructure that can meet the needs of any customer...

That knowledge needs to be immediately available to the next patient... who might log on in China. It's another level of real-time processing.

COOK Report: As the engineers like to say, that is not trivial.

Zakim: I think it's very hard to build an infrastructure that absolutely meets the needs of *any* comer. But it appears the Dutch have probably built an infrastructure that makes it relatively simple to meet the needs of a variety of applications.

COOK Report: Bob Hertzberger's whole approach to building these e-science stacks has been to make them as modular and generic as he can. The goal is that when various disciplines come in and want to add their own tools and their own approaches, that they can do it as cost-effectively and time-effectively as possible.

Zakim: Even if they haven't gotten there yet, these are still very interesting people to talk to. These are the people who want to make it possible.

How to Organize the Engine to Let Innovation Flourish?

COOK Report: Kees Neggers and Bob Hertzberger are amongst the leaders in the Netherlands articulating the research philosophy that David Zakim found so fascinating when I used Kees'

NORDUnet articulation that follows.

As one looks at the ongoing need to fit the innovative Dutch patterns to the larger and more complex technology plans and roadmaps of the European Union as a whole, one begins to bump into different planning philosophies. What follows is a narrative summary of the approach that Kees Neggers outlined at a NORDUnet meeting in September 2009. I let Kees paint his vision and then add my own commentary. The Powerpoint file and video of the presentation are available on the NORDUnet website under speakers data: http://www.nordu.net/conference/ndn2009web/programme.html

Neggers: I want to explain my perspective on the approach to networking and knowledge infrastructure for research that we are taking in the Netherlands now.

There are three trends in research...

First there is the emergence of System level science which is defined as "the integration of diverse sources of knowledge about the constituent parts of a complex system with the goal of obtaining an understanding of the system's properties as a whole" [Ian Foster].

Next is Multidisciplinary research. Here each discipline can solve only part of a problem. Consequently you need collaboration between different research groups. These groups may be distributed across states, countries, or continents.

You have Research driven by (distributed) data. This happens in the context of a data explosion, both in volume and complexity. You want both simulation and experiment combined. Your scientists want to be able to explore data-sets with no up-front hypothesis.

All of this requires cutting edge network innovation for which internet is inadequate. For these needs a "Best effort" network is not good enough. The Internet, as a best effort network was designed for delay tolerant, many-to-many communication.

…a "best effort" network is not good enough…

To meet our research needs, the network will have to deliver guaranteed performance for large data flows and time-critical applications and do so across multiple domains with different technology while keeping the successful end-to-end model of the internet.

Other changes are leading to new requirements for networks as infrastructure. We have an explosion in the amount of data from experiments and simulations. Here are examples: LHC, LOFAR, e-VLBI, ITER, Climate modeling, Humanities and Social Sciences. We also need near real-time processing of very large datasets. It is necessary to be able to support an increase in remote collaboration associated with distributed sensors, shared computing and storage, grids and finally Virtual teams.

We need a seamless, integrated ICT infrastructure. An integrated ICT infrastructure must facilitate the shared use of networks, computing power, data storage, instruments, etc. It must ensure hassle free end-to-end connectivity. It must be there as an infrastructure into which like an electric utility users can plug. It must provide a single user interface and a single control plane for the allocation of multiple resources, from multiple domains and in multiple locations. It must provide an integrated set of services to support all kinds of research.

Necessary for the Future - A National Integrated Research Infrastructure

Our goal is to create an integrated ICT infrastructure for the Netherlands. In June2008 SURF made a proposal to the Dutch government to integrate existing organizations for Research Network, Grid Computing, Supercomputing, and services for e-science into a single organization responsible for the ICT infrastructure. And to finance the development of a single control plane and a consistent set of services for research on a structural basis.

In December2008 the Netherlands ICT Research and Innovation Authority (ICTRegie) provided a similar recommendation with SURF as umbrella organization. In May 2009 the Dutch government agreed with the principles and asked SURF to work out an implementation plan.

Implications for governance models

Governance models for the ICT infrastructure need to encourage innovation in order to be able to offer advanced services without sacrificing reliability. We also must ascertain how we can achieve all this. In my opinion this should be done by coordination and collaboration, in close relation with advanced users, not by creating monopolies

Bottom up innovation

- Innovation cannot be planned, designed and optimized
- Innovation is like evolution: unpredictable, but with a clear direction (at least in hindsight)
- Evolution needs three elements:
1. Variation: many attempts to solve the same problem
2. Selection: an environment where the best solution is given a chance to thrive
3. Propagation: dissemination and replication of the best solutions [Source: Prof. Bart Nooteboom]

The choice appears to be either a Centralized governance model where on the apparent plus side centralized "command and control" over the entire infrastructure may avoid duplication of efforts and introduce economy of scale.

Innovation is like evolution: unpredictable, but with a clear direction

On the other hand one loses variation and central planning tends to stifle innovation. In addition, in large communities decision making will be slow and will lead to mediocrity.

With a Decentralized governance model you get decentralized "unity in diversity" and the ability to experiment with multiple solutions in parallel. You get closer to user needs and can facilitate organic growth. The price will be some degree of duplication fostering a need for coordination.

How to decide on the best model...

A centralized model is likely best for large-scale, single facilities (e.g. LHC). It also works for distributed, homogeneous sets of resources (e.g. LOFAR). On the other hand we think that the decentralized model is a good option for distributed, heterogeneous facilities (e.g. EGI, CLARIN, LIFEWATCH). The best example is the internet where you have some coordination of communication standards and the use of identifiers while everything else is decentralized.

Is cost efficiency a valid argument?

A centralized model may introduce economy of scale, but the costs of research networking are not determined by the international backbone network alone, but much more by the NREN and campus networks. Moreover, economy of scale of a central approach is minimal com-

pared to the spill-over effects of innovative services. Hence, a centralized model works well at EU(or larger) level for very expensive one-off expenditures for large facilities like LHC or ITER, not for pan- European communications infrastructures If cost efficiency were the argument, we should all be using the commercial operators!

Why network innovation needs a decentralized approach

Central organizations tend to become static. Why? More parties are involved in synchronous decision making at all levels leading to more talking and less action. Effort must be expended to hammer people into a single organization, single vision and a single approach.

...technology organization works most effectively when it percolates bottom up...

In contrast decentralized collaboration tends to create innovation. Smaller entities can offer more flexible and quicker response to a changing environment. More parties working in parallel create different visions and "friendly

competition" of ideas within an agreed open interworking environment.

COOK Report: As John Seely Brown and John Hagel say in their collaborative blog *Edge Perspectives in the 21st Century* and in the internet the only locus for creativity is at the independent edge. The center or the network exits to provide a common protocol by which the edges may communicate with each other.

It is here that the European Union formed to make a global economic power out of the smaller member states in the age of internet linked in-stantaneous communication and collaboration fails to de-liver value. By lowering eco-nomic and trade barrier be-tween member states, it can add value from the existence of collective whole to the in-dividual members.. However as far as technology projects go, given the fundamentals of the network and how it enables the simultaneous de-centralized evolution of col-laboration, the time taken by getting the edges to agree to take directions from a pan European, center-directed technology plan will mean that the edges have lost from one to three years of time in which they could have been doing things on their own.

Furthermore because the hardware delivers ever more bang for the same buck and the e-science center has been designed to commoditize the production of open source e-science middleware, Brussels has in this area only dimin-ishing returns to offer if the project does not rest on some huge capital expense like the LHC. There is a say-ing that all politics is local. To it could be added the corol-lary that technology organi-zation works most effectively when it percolates bottom up from the plans of people who will pay for and use it.

VIII. How A Progressive ICT Infrastructure Benefits the Economy

The innovation engine as open fabric

Through a number of advanced ICT infrastructure projects, the Netherlands has become an attractive country for scientific research. The hub function of the Amsterdam Internet Exchange (AMS-IX), the global connectivity achieved through SURFnet, and the active participation by Dutch groups in various international projects all have all contributed in attracting high-tech companies and research centers. This has led to a rich ecosystem of ICT related activities, as well as research activities enabled through the ICT infrastructure. Examples of this type of concentrated activity can be seen in the High Tech Campus in Eindhoven and in the Amsterdam region.

The concept of e-science as depicted in the science cases shown in the preceding chapter is innovative by itself: completely new forms of research are possible with the availability of an ICT research infrastructure that provides state-of-art computing, storage and communications, as well as instrumentation for the researcher. In this section we focus on the innovative aspects of the infrastructure on industrial R&D and as a stimulus to economic activity because the contribution to academic research has been sufficiently highlighted under Science Case.

One of the innovative aspects of an advanced ICT research infrastructure is the potential to perform cutting edge research in a number of fields and consequently reducing R&D lead time. Through collaborative projects with industry by academic research groups, an automatic transfer of knowledge is realized. This mechanism and the results from VL-e, MultimediaN, NBIC, BiG Grid and GigaPort have already led to extensive collaborative efforts with major industrial partners including Philips Research, Organon (then part of Schering-Plough now part of Merck), DSM, and Logica/CMG.

> *…new forms of innovation become possible with state-of-art computing, storage, communications and instruments…*

An important innovation driven by e-science methods is a change in the research method itself. The infrastructure becomes an instrument, delivering the flexibility to set up different partnerships as and when required.

Philips provides a good example. In recent years Philips Research has undergone a number of major changes, refocusing from chips and consumer products to health systems. Philips Research now finds itself with a different set of priorities and needs to work on subjects such as biomarker discovery where there is far less expertise than in the traditional areas.

To accommodate this focus shift, Philips Research has adopted the open innovation concept, which has resulted in a more open organization, (especially for a company founded in 1891.) It stimulates strategic cooperation with other companies, universities and research institutes. Facilities are shared with partners in order to improve cost-effectiveness and efficiency. The Research ICT Department has modified its strategy in order to enjoy full advantage from this change.

The effects of innovation in GigaPort NG Network
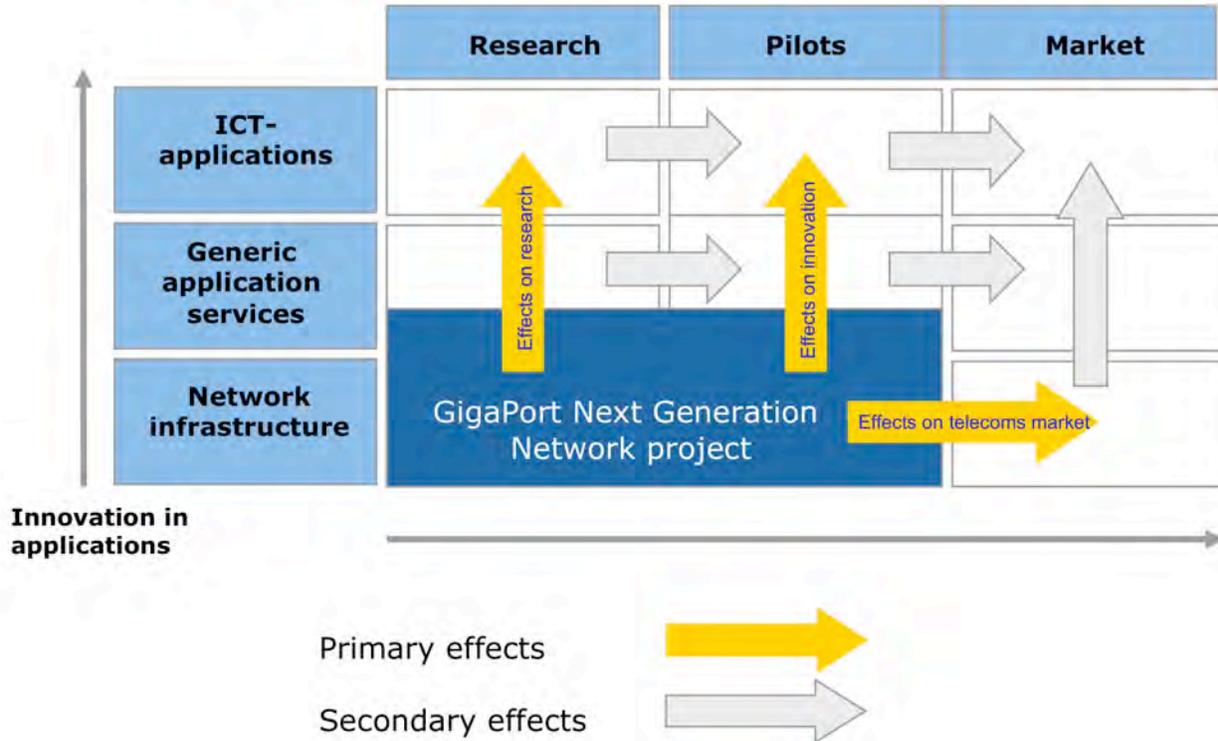
Research networks and e-Infrastructure for Research in general act as innovation agents. The need to serve a demanding user community of leading edge researchers forces e-Infrastructure providers to invest in pilot services ahead of the market. In doing so they create an innovation engine that works in many dimensions: it creates a demand pull for new networking research and products and a market push for new telecommunications services. At the same time it triggers development of new middleware and applications, which then themselves function as engines for innovation in their layers. As a result it not only shortens the development cycle for new services, in many cases it allows services to mature which would never be able to do so on a purely commercial basis as is nicely explained in the interview with Hans Dijkman later in this report.

Through collaboration with GigaPort, VL-e, and BiG Grid (for which Philips Research will host a significant part of the infrastructure), the department is now able to provide networking and computing facilities to the research groups - an action on the part of Philips Research which would not have been feasible in any other way.

While the 'Philips' case is already very convincing, more industrial involvement is on the way. In the NBIC consortium several industrial partners (e.g. Keygene) have also expressed their intention to become partner in the distributed ICT research infrastructure, starting with housing a 'life science cluster'. This will closely follow the model chosen for Philips Re-

search. Similarly, Organon has moved its main focus from bio-informatics to e-Bioscience, allowing it to adapt quickly to new topics in drug design.

Building an ICT infrastructure for research creates spin-off for industry

Not only does the ICT re-

search infrastructure enable innovative research in the larger industrial community. Building and evolving such an infrastructure itself is a stimulating formidable challenge in itself. It creates many opportunities for innovative activities. This will undoubtedly stimulate regional and national economic activity.

A parallel can be seen with the development of the Internet and the World Wide Web, which – in the Netherlands – started largely in and around scientific institutions. This has led to the emergence of a great variety of ICT related business activities – as has been demonstrated in other countries as well. The ISP industry (Internet Service Providers) grew out of the slipstream of the development of the Internet, which has been boosted by SURFnet, it is more than likely that the development of generic services for research will stimulate the emergence of services-based industries, offering a new generation of opportunities and economic models.

Research networks and e-Infrastructure for Research often serve as innovation agents. The need to serve a demanding and knowledgeable user community of leading edge researchers forces e-Infrastructure providers to invest in pilot services and to

pioneer better customer support strategies far ahead of the market. In doing so they create a fabric-like innovation engine that works in many dimensions: it creates a demand pull for new networking research and products and a market push for new telecommunications services. And, as has happened here, it triggers development of new middleware and applications, which then themselves function as engines for innovation in their layers. This shortens the development cycle for new services, and in many cases it allows services to mature that would never have been able to survive in purely commercial marketplaces. (This process is defined and explained more fully in a conversation with Hans Dijkman in the next chapter, "SURF as Economic Midwife for Technology Transfer".)

Innovation is contagious

Both the research institutes connected to the infrastructure and the vendors supplying the various components to ICT researchers and support teams are themselves often inspired to innovate within their own environment. SURFnet is a good example. In the GigaPort and GigaPort Next Generation projects, SURFnet has challenged both vendors and us-

ers to push the envelope in advanced networking. Organizations connected to SURFnet have been encouraged to upgrade their own internal infrastructures to make the best possible use of available services. Vendors were encouraged to implement the newest technology in their offerings.

Service providers have been able to use SURFnet as a platform to experiment with innovative services for sophisticated user communities and early adopters, before rolling out to the wider population. As a result, not only the research community but the ICT community as a whole has benefited from the innovation initiated in the GigaPort projects.

The picture on page 71 above illustrates the 'innovation engine' effects of GigaPort.

Part of this is a general enthusiasm for progressive use of new technologies. There is also an almost patriotic recognition that this is supposed to grow the economy and improve the quality of business offerings. It's in the unwritten job description that everybody involved with evolving technology is also involved in new businesses, new products, new employment and new revenue.

At it's most productive, the productivity and value-
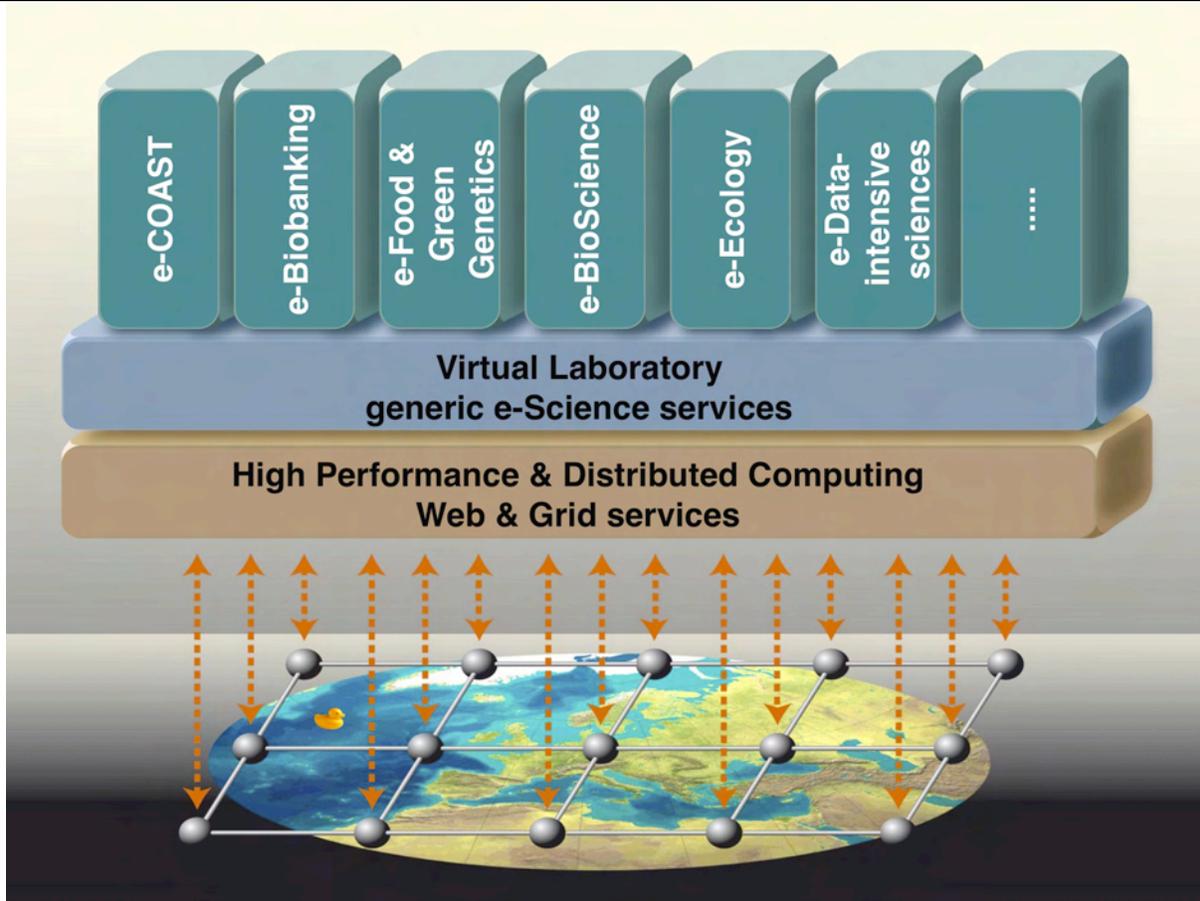
An illustration of the Netherlands "e-infrastructure" from Bob Hertzberger's VL-e project.
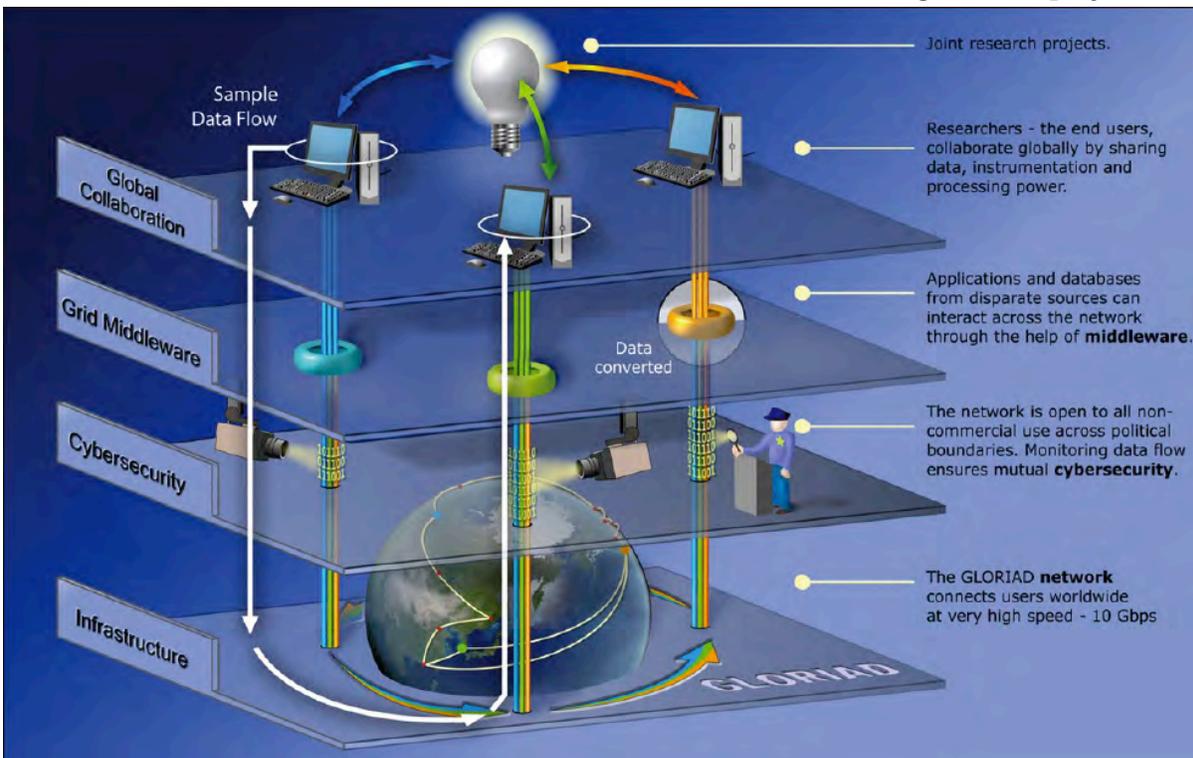
**A diagramatic view of what in the United States we call "cyber infrastructure".
The VL-e picture shows the multidisciplinary collaborative use of the infrastructure, while
GLORIAD picture shows the data flow resulting from worldwide collaboration.
For more information see the Gloriad web site http://www.gloriad.org/gloriaddrupal/**

generation that Kees Neggers and Bob Hertzberger have described on the pages of this document is what Carlota Perez would call a case of radical centralism. In the words of Harvard Professor and sardonic songster Tom Lehrer, (a favorite of a earlier generation of geeks), it's about "doing well by doing good."

Partnership Case

By definition, an ICT Research Infrastructure on a national level is a shared concept. This does not only hold for the collection of hardware resources and generic software development, it also holds for the individual resources themselves: the network, computing systems and storage facilities. For years, scientific communities in the Netherlands have collaborated in realizing large-scale facilities as a superior research network and large high performance computing systems. Collaboration in this way has not only led to economies of scale, but has also led to new styles of multidisciplinary scientific work, as a result of using large ICT research facilities together. Furthermore, it appears that in The Netherlands, both academic and industry research groups are involved in this type of collaboration.

SURFnet has challenged vendors and users to push the envelope

Many scientific communities in the Netherlands are using these national infrastructure facilities as a basis for their collaboration. For example, the SURFnet6 network was created by a broad consortium of academic and industry partners, for use by the entire research community. The ICT infrastructure provided through BiG Grid will be used by the Nikhef high energy physics community (linked to LHC at CERN), by the astronomy community (LOFAR), by the life sciences community as represented by NBIC, by alpha science communities like DANS, and by computational sciences communities represented by NCF."

an almost patriotic recognition that this is supposed to grow the economy…

Due to its active attitude towards international cooperation, and to its advanced infrastructure, the Netherlands is in a unique position to attract international collaborations as well. This has already led to significant Dutch

presence in networking organizations as GEANT, GLIF and TERENA, in HPC-oriented communities such as DEISA and PRACE, and in grid-oriented communities such as EGEE and EGI. [**Editor's Note:** Cees de Laat noted that significant leadership roles and contributions are also done to IETF, OGF and W3C. And Bob Hertzberger as another illustration adds that in mid 2009 Nikhef won the bid for the European Grid Infrastructure (EGI) head office.**]** The innovative lightpath exchange NetherLight has grown into an important hub for networking connections, leading to new partnerships both at the network and the applications level.

Technical case

General description
The science and innovation cases described previously require networks, computing, and storage technology. Linking these together in a coherent and more seamless way, using web and grid technology and employing e-science methodologies, is a *condition sine qua non* for progress in almost all scientific fields. Thus a landscape of resources must be formed, ranging from networking, capability and capacity based computing resources (see under computing and storage resources), storage and data services, to application support services.

…cutting edge high performance networks are indispensable for modern research…

The total chain of resources, middleware and services determines the added value of the infrastructure for the researchers. The people who become skilled in working in such an environment are a resource in their own right. The human actors in a progressive infrastructure can be recognized as becoming invaluable "components" of the infrastructure as well. It is often these infrastructural actors who choose to break off and start a new business or create a new innovation community.

This infrastructure can only be realized by developing the necessary services through an integrated effort – and with a commitment to open technology. This is what assures opportunity to re-use the data processing services and thereby the re-use of knowledge. More value is thus created.

A layered approach

Schematically, the infrastructure model is represented by the figure below. The technical case addresses the ICT research infrastructure, as

contained in the highlighted box, which ranges from the network to generic e-science services and the tools to deliver them.

The Network

The availability of cutting-edge high performance networks is indispensable to modern research, as it connects distributed scientists, instruments, computing and storage facilities. These networks are no longer distinct from the rest of the ICT infrastructure: resources within the network and within the rest of the infrastructure will have to be controlled through an integrated middleware layer, to ensure an optimal allocation of resources.

ICTRegie facilitates the formation of innovation platforms

SURFnet6, the current generation of the Dutch national research network, is a state-of-the-art network and is widely recognized as a world leading network. SURFnet6 is a hybrid network, providing both IP connectivity and lightpaths within the same network. Lightpaths are dedicated high-capacity, low latency point-to-point connections between two nodes in

the network. These properties make them very well suited for transport of large quantities of data, such as originates from scientific experiments, or for high performance interconnection of grid storage and computing components. snip

Computing and Storage Resources

Capability-based Computing

Different problems have different computing needs. Allocating the most efficient computational resources to problem areas not only has scientific advantages (only those groups in need of certain high end high cost hardware will gain access, thereby increasing efficiency), but also has budget advantages. This has long been recognized in The Netherlands, leading to so - called "capability resources" and "capacity resources".

ICTRegie and the Dutch ICT Infrastructure

ICTRegie, the Netherlands ICT Research and Innovation Authority, as its web pages explain: "is a compact, independent organization consisting of a supervisory board, an advisory council, a director and an office. The Ministers of Economic Affairs, and

of Education, Culture and Science bear the political responsibility for ICTRegie. The organization is supported by the Netherlands Organization for Scientific Research (NWO) and SenterNovem.

ICTRegie effectively gathers together the stake holders for ICT. It engages in demand articulation which means stimulating and keeping track of (latent) demands from customers and businesses. ICTRegie encourages this activity by facilitating the formation of ICT Innovation platforms for different areas of expertise (e.g. Health Support, Product Software). Within these Innovation plat-forms, parties like knowledge institutes, ICT businesses, and representatives of potential users or customers come together to explore meaningful utilization of ICT and collectively formulate topics for further research." http://www.ictregie.nl/ictregi ehome-FAQs.html .

Above left: Kees Neggers, Cees de Laat and Gordon Cook and above right Bob Hertzberger, Hans Dijkman, Cees de Laat and Gordon Cook at Supercomputing 09 in Portland Oregon.

IX. SURF as an Economic "Midwife" for Technology Transfer

An interview with Hans Dijkman, Kees Neggers and Bob Hertzberger

Neggers: Meanwhile Hans Dijkman has some slides explaining how the SURF innovation model fits into the larger economy as a whole.

Dijkman: Let me talk about what I will call "Product Development 101." Over time research is done with the aim of introducing new products.

The functional composition diagram below illustrates the e-science infrastructure that produces the research that we are talking about now.

We are putting an e-science Research Center between the 'hard' infrastructure and the specific science domains. The center is, for example, responsible for making generic e-science tools or middleware that can be used by the specific scientific domains. The activity is not limited to only engineering (i.e. a form of applied science) but also research is done to develop new scientific methods and techniques that are made possible with modern ICT.
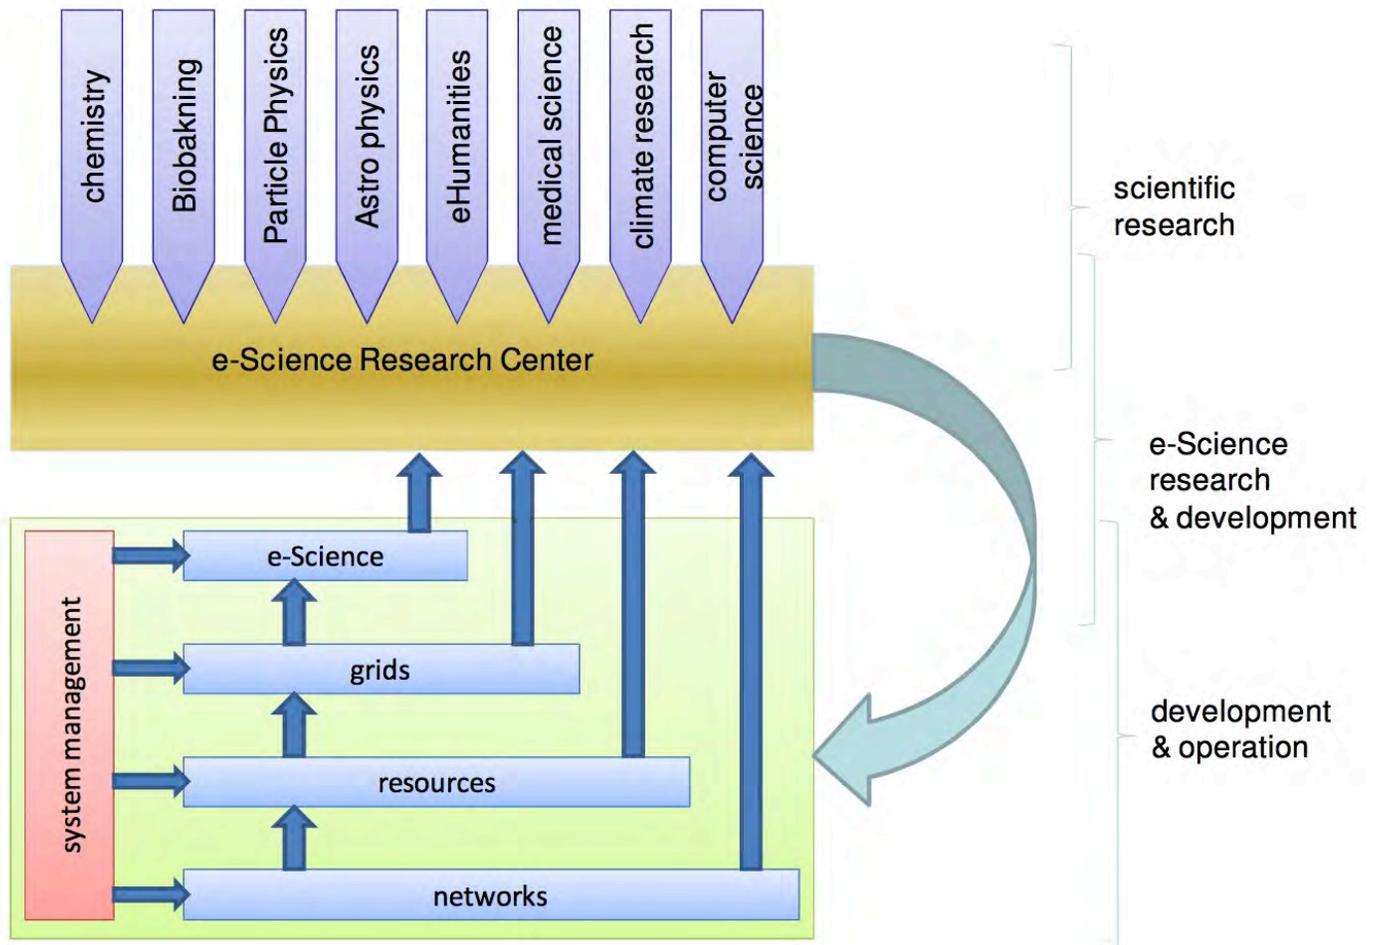
Functional composition of the E-science Research Center and ICT infrastructure for scientific research. Source is ICT Regie.
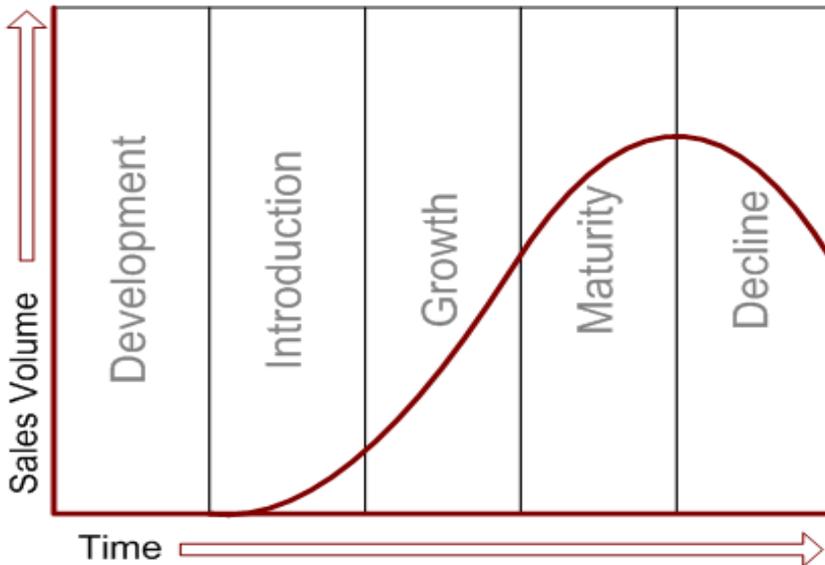
The classical Product Life Cycle. The different stages of the 'life' of a product in time.

which we call the starting point. Here we have fundamental research and it is unknown whether there will be market pull or market push. But over time some demand can emerge. People see it and say- "oh I could use that." They become the first users from the general domain in which the research is done. This is the first phase of demand. When you have enough demand at vertical axis A, supply begins to be offered.

At vertical axis C, supply and demand meet in a plateau and a competitive market for the new product emerges.

On the left side – axis A - the product is purely new and innovative, but in the middle as other people begin to want it, you have a non-competitive market beginning to develop. At vertical axis B, standards are beginning to

Applied and fundamental sciences interact with each other in the center. In our vision e-science stands for 'enhanced' sciences. That also implies that the interaction is collaborative. Different domains can share each other's resources independent of time and place.

For financing we used a special approach. To illustrate this we use an expanded version of the familiar Product Life Cycle (PLC) model to explain the need for an e-Science Research Center and the need for governmental funding.

We all are familiar with the theory of the PLC, usually presented as the following This model does not represent the product development stage or the preceeding innovation stage. To describe

the innovation stage we expanded the model.

…three market stages:
basic research…
non-competitive market…
and competitive market…

For simplicity we assume that every product can be traced back to some basic research
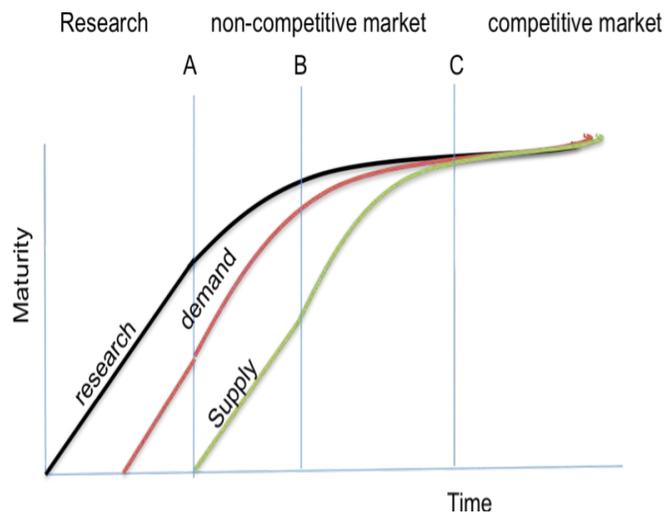
The 'life' of a product in terms of its maturity in time. From 'birth' to a standardized commodity. The demand and supply curves are –delayed- following the product life.
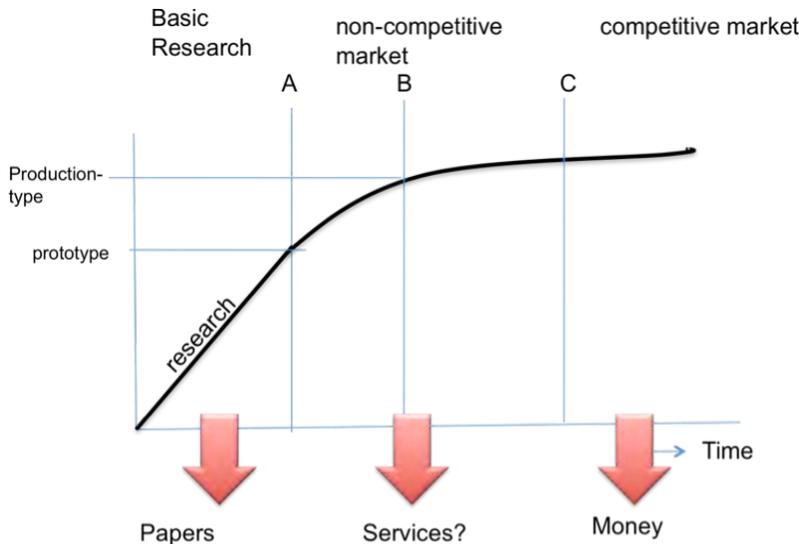
The 'life' of a product in terms of its maturity in time. But now added the 'drivers' of the different stages.

companies, for example Philips and AT&T, are able to cover the cycle from basic research to the market. For them this middle stage is used e.g. to obtain patents and the IPR. In this way they get a competitive advantage. Let's call this Proprietary Innovation. Other companies are starting from existing technologies, in other words, when the academic R&D (sub)stage is finished. In the case of e-science, the goal is to develop new methods and techniques by which we do science. E-science leads to Open Innovation. This means that, under certain conditions, ideas, results and products are shared. The financing should reflect this character.
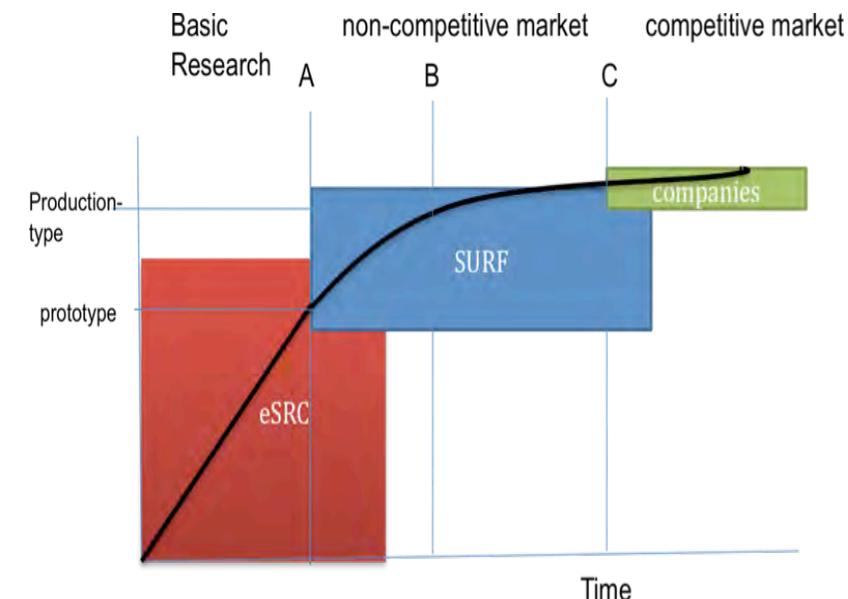
The output/results of basic research is rarely ready to market. The quality requirements are completely different. So, in this stage, the re-

emerge and this emergence enables the increase in supply to develop into a competitive market. Altogether, we distinguish three stages; basic research, non-competitive market and the competitive market.

The output of basic research is rarely ready to market.

What are the driving forces in the different stages? Basic research produces publications which are what keeps people at universities alive. In the competitive market is it simple: money. But what about the middle stage?
In industry the term "research" is frequently used to describe innovation with ex-

isting technology, which academic scientists would normally describe as development. This different use of the word "research" can lead to many misunderstandings. If we represent the middle stage as follows, [see non-competitive market time chart below] the problem becomes more clear. Only a few

The non-competitive market stage of a product but now in more detail.
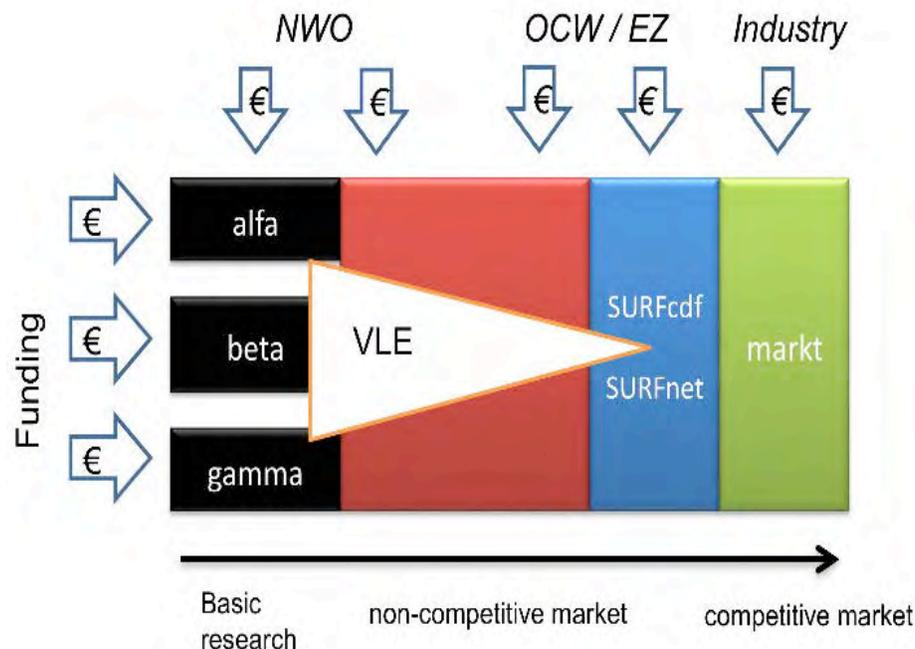
sults from basic research have to be upgraded to a level so that others can use them. Prototypes have to be engineered into production types, computer applications have to be made more user friendly, and so on. An important part is also standardization. For e-science, it is the stage where applications are made generic and methods are developed to give access to different resources.

It is a necessary stage but with no clear incentive or result. It's not where you spent your basic research money. It is also not where companies invest for their product development. This "in between kind" of stage needs an external force, a push or a pull. At this point Venture Capitalists show up in case of potential commercial value. In the case of e-science this is not working. It is here where additional governmental sponsorship is needed. Otherwise nothing will develop.

Different driving forces require different organizational models. Therefore it is nearly impossible to put them under one management. If we focus on the e-science research, we have made a choice for different organizations for different stages.

So SURF here in the blue section immediately below provides the bridge of support between the early adopters and the arrival of a competitive market. The early adopters are the pioneer users you need in order to mature your products and services.

You make something. Someone likes it and starts to use it and then you get the demand for help. Users want a help desk. Using your grid library needs other forms of support. Until it becomes a commodity that you can buy in the marketplace, you need

Regardless of the nature of the product, humanities, behavioral sciences or science, the maturity is about the same. In the e-Science Research Center they can come together as building blocks for an e-science application. It also shows in which stages VL-e was operating.

a source of support for this. But also engineering is needed to transform e.g. an application from prototype into a more robust, user friendly, etc application.

*This is the glue the government has to inject…
to shift from research to market…*

Above is where the different scientists are funded by the Dutch science foundation – the NWO. The three research domains are alpha – humanities, beta - hard sciences, and gamma - behavioral sciences.

COOK Report: The above VLE triangle is the development investment flow?

Neggers: The important thing to note here is that the red blue and green are normally vertical pillars doing things entirely on their own. Through VL-e (Virtual Laboratory for e-Science) the black boxes representing the three broad types of research cut across the otherwise independent agencies and make interdisciplinary research possible.   There are some tools that you need within single research domains but other tools are needed as well that cross scientific domains.

COOK Report: OCW is?

Neggers: The Ministry of Education, Culture and Science.   This Ministry gives money to SURF and NWO – not the ministry decides which researchers get what money.   NWO is going to finance this multidisciplinary e-Science Research Center (eSRC).   And based on the ICTRegie advice the ministry is assumed to fund the blue pilot phase under the SURF umbrella via SURFnet and SURFcdf. SURFcdf   is the working title of the still to be financed High Performance Computing and Data Facility.

SURFcdf is a work in progress.   The intention is that the money spent on supercomputers and on grid and data storage projects are merged into a single facility making it available to all disciplines.

COOK Report: So this is how you have institutionalized the support necessary to bring basic research to market?

Neggers: Bob, correct me if I am wrong but your VLE project was successful and now ICTRegie has advised the government that when this project ends it should become a new organization not just for the few application people that you were able to serve with your tiny bit of money, but it should be

larger and permanent – not just a project.

Hertzberger: Correct.

Neggers: This is the glue that the government has to inject to make this shift from research to the market place. While it also stimulates multidisciplinary applications for the first time.

Hertzberger: You can never get these guys to do this themselves.   There is no incentive and they lack the experience.   You need multidisciplinary interaction between computer scientists and those domain scientists. You need it but don't have it and consequently you setup an accelerator to leverage existing applications and develop new kinds of applications.

Neggers: But you also need to underwrite the computer science part of the entire middleware that is needed to permit progress along this path to establishing this new paradigm for interdisciplinary e-science learning.   It has both a computer science and an application component. It is the merger of the two that creates the added value by leaving the generic part to the computer scientist and the application part to the domain scientist. Together they have to make something that works.   You need them both.

Dijkman: Correct and if you do it right you make way for what Henry William Chesbrough was talking about in his 2003 book Open Innovation: The new imperative for creating and profiting from technology. I quote:

"In today's information-rich environment, companies can no longer afford to rely entirely on their own ideas to advance their business, nor can they restrict their innovations to a single path to market. As a result, says Harvard Business School professor Henry W. Chesbrough, the traditional model for innovation - which has been largely internally focused, closed off from outside ideas and technologies - is becoming obsolete. Emerging in its place is a new paradigm, 'open innovation', which strategically leverages internal and external sources of ideas and takes them to market through multiple paths. This path-breaking analysis is based on extensive field research, academic study, and the author's own longtime experience working in Silicon Valley. [The book offers] rich descriptions of the innovation processes of Xerox, IBM, Lucent, Intel, Merck, and Millennium, and the many spin-offs that have emerged from these firms. Open Innovation shows how companies can use their business model to identify a more enlightened role for R&D in a world of abundant information, better manage and access intellectual property, advance their current business, and grow their future business. Arguing that companies in all industries must transform the way they commercialize knowledge, Chesbrough convincingly shows how open innovation can unlock the latent economic value in a company's ideas and technologies."

e-science needs to happen in an open environment...

I have talked about 'open innovation' versus proprietary innovation. I believe that e-science must happen in an open environment. No IPR, no patents, etc. I also think that the outcome of e-science, e.g. generic tools to support applications should be handled like Open Source Software.

Most Webservers are based on open source software called 'Apache'. The core of Apache is maintained and managed by the Apache foundation. They protect the quality, etc. Everybody else is not allowed to change the core but is allowed to build add-ons. There are more examples.   You can imagine that something equivalent to the Apache foundation is formed. This 'foundation' develops and maintains a set of e-science applications, etc. But e-science is also sharing of e.g. databases. Between academic institutions no problem but sharing data between companies becomes a lawyer's paradise. A spin-off of e-science should be, in my vision, a kind of Clearing house. A place where data can be shared without the legal bureaucracy.

X. Coming to Conclusions

Realizing the benefits of a re-usable infrastructure

COOK Report: Lessons learned in the Netherlands should be an education for the rest of the world as well. It is a testament to the power of the commons — and to the willingness to learn how to work together. It is about the benefits the Netherlands has demonstrated can result from sharing the network, the knowledge, and the tools. And it is that "spill-over" effect Kees Neggers of SURFnet and Wim Liebrand of SURF keep mentioning — in which the benefits of creating a national knowledge infrastructure are widely distributed and pay back results which of exceed expections.

At the same time, we have in comparison the rest of the world — the United States included – where so many have chosen to focus on private not public good. Where the consequence has been one of short term horizons rather than long term investment. That needs to be changed, or at the very least moderated.

The power of networked technologies to change the rules going forward cannot be denied. As the physical world runs short of opportunities, the digital world offers new frontiers. The future will belong to those countries and communities who are proactive in investing in national infrastructures that serve the many rather than the few. Over time the world will see that infrastructures left in walled gardens will wither.

Lessons learned in the Netherlands should be an education for the world

An area where it is easy to recognize the growing opportunity enabled by open infrastructures is e-science. With e-science it is "early days". The temptation is to compare it to where we were in 2000-2001 with blogs and wiki's. People were experimenting with both, but neither the tools nor the culture were in place to enable really systemic widespread use. We needed another six years before the possibility for global collaboration was manifest and could be articulated by Yochai Benkler in his book The Wealth of Networks. http://en.wikipedia.org/wiki/Yochai_Benkler.

The earliest perception of blogs as reported in the media – that they were the mindless daily diaries of millions was faulty. Of course there was some truth to this but what really mattered was that blogs enabled the forming of community of interest and joint problem solving. In effect, they began to "grow" fresh new commons that increasingly inter-connected. And that takes us back to e-science in the Netherlands.

As discussed earlier in Chapter V "Making e-Science Work" there are at least two vectors of differentiation for e-science being pioneered by important e-science initiatives like VL-e.

The first is supporting a methodology for science that embraces trans-sectoral and

with limited funds decision-makers had better put it in the infrastructure first

multi-disciplinary collaboration. The second is creating generic middleware tools that suddenly make it possible for

different disciplines and communities to begin collaborating.

The future will belong to countries and communities most proactive in investing in national infrastructures.

Science in the past rewarded the expertise of the "hedgehog" who was totally focused on a narrow specialty and encouraged to stay in their silo. Today in a world of growing ubiquity of connectivity, the hedgehog will find himself useless, unless he can figure out how to weave himself into a larger fabric. Increasing, as Cees de Laat said earlier in this volume, science is now global. To get at real problems and generate real results, science must go where-ever the knowledge trail takes it - must cross geographic borders, and knowledge domains as well. What does that require in terms of infrastructure.

SURFnet laid the basic utility grid and the basic foundation on which everything else is built. Then the team considered basic progress in optical networks and understood the cost effectiveness of doing as

much as possible at layer 1 and layer 2 and going to three only when absolutely necessary. It collaborated with the likes Steve Goldstein and Tom DeFanti and started slinging 10 Gb lightwaves from Amsterdam to Chicago. This was in 2002-2003. But immediately Cees de Laat found it necessary to adjust the behavior of TCP/IP in 2002-2003. (Doing real infrastructure is not for wimps.) That allowed SURFnet to do things demonstrate the ability to conduct real-time four times high-definition telepresence sessions in Japan and North America and Europe.

….adjust behavior of TCP/IP to optical network properties.

Now, at the same time, and in parallel with its FES infrastructure calls, the Netherlands team is working on defining what the next level will be. It has pioneered the next highly important integrative stab of rethinking the network stack up to the application layer. The application layer is the rubber meets the road for e-science users. It's where user connectable lightpaths and Web services, remote instrumentation and grid and preliminary e-science middleware are be-

ginning to enabling the equivalent of the blog and wiki for global science. And as the Fourth Paradigm book has shown, there is already tremendously impressive work going on globally in data intensive science.

But this must be done in parallel. Technology never sleeps. Optical development hasn't stopped in Nortel's common photonic layer. It keeps on getting better. Anyway, bringing all this together in the application layer has stimulated the beginnings of a truly global user science movement where people in quite a few other countries besides the Netherlands are beginning to experiment.

In the process of preparing this document, Bob Hertzberger, Director of the Virtual Laboratory for e-Science, wrote the COOK Report that "for e-science to work a huge effort has to be put into innovation. But technology developments determine the pay back to customers." This stimulated additional conversation with Hertzberger.

Hertzberger: We have to put tremendous effort into innovation before we can illustrate that it pays back to our customers. How long, and more important, at what rate we absolutely do not know, because it depends on

the speed with what technology around you develops. The only thing you know for sure is that when you don't do it you will loose.

For e-science to work a huge effort has to be put in innovation, but technology developments decide the pay back to customers.

Consequently what you do is try to take the competitive advantage. We have had some visionaries (especially Hans Rosenberg) who saw early on the importance of networking for a small highly developed country. The e-science case has been made possible by the existence of an excellent network and the fact that science, in particular medical science, is well developed in the Netherlands.

But most important of all that, as a young scientist, I had the privilege of working in an experiment which got the attention of the whole world (even the New York times) and for which the leader got the Nobel prize [the discovery of the W and Z bosons in high energy physics]. From that experiment I learned the importance of

instrumentation as well as the fact that you have to invest a large number of years before you get your return on that investment.

However as politicians are not scientist or entrepreneurs you seem to have to promise that the sky is the limit before they believe (and please notice I don't say understand) that what we are doing has some relevance and must to be funded.

COOK Report: This is called building infrastructure and you do not get payback from infrastructure in 90 days and maybe not even in 900 days but you do get incremental payback during the life of the infrastructure.

Science globalization demands for education towards multi-disciplinary scientists

Hertzberger: Yes and note again a scientist is educated to solve a problem in his often mono disciplinary field. However, at the same time, because of the potential of networking, science is also becoming a global activity. That consequently e-science has to support that globalization first before the end user can obtain full advantage of e-science for his applications.

Like it or not this requires changes in the sociology of science. Specialized young scientists had better understand that e-science - as a form of instrumentation - will irreversibly change their profession, and that because science is conservative they better start today rather than tomorrow. That in order to realize its advantages they have to accept globalization in their domain and that a consequence is that they have to work in more multi-disciplinary teams. They have to realize that the data/information centric angle – being collaborative - is often considerably different from the way they have learned to operate up till now.

COOK Report: But your e-science infrastructure building process needs money to keep it going. Consequently it is necessary to explain to the funding agencies how it works and how it relates to developments in other countries in the world and how one may most accurately think about the dividends that one can expect to see. But such dividends that cannot be mature in 90 days or 180 days without the basic foundations needed for global collaborative scientific inquiry and even more general learning

Hertzberger: I would say it this way. That for the gov-

ernment and the decision makers the problem is that they do not have unlimited funding and that consequently they had better first put it in the infrastructure. However, because this is a new game, established scientific communities all will claim they can do it far better in their field of science if they don't have to spend time on all these new things because it only slows them down. There is a continuing tension between the old ways and the new. Consequently there is a huge risk that every field is re-discovering the wheel over and over again and that you have to prevent that by adopting generic solutions and re-use of components.

…mono-disciplinary attitude results in huge risk in rediscovering the wheel over and over again…

COOK Report: And the point is, as the Netherlands has shown time and time again, if you don't do it first others will. There is a prime mover advantage. It is definitely better to be building the third or fourth story of your building while your neighbors are only completing the second floor. And if you use generic components like you used the saw mill in the 17ᵗʰ century you can build new floors on your edifice faster than the others.

Hertzberger: Yes. in a modest way building up a reusable infrastructure is exactly what we try to do in the VL-e project.   And then I have the vision that similar ideas like the ISO/OSI model which was designed as a reference model for networking and the virtual machine model that has been used to avoid that for each new microprocessor you had to design a complete new compiler chain might also work for part of that e-science infra-

structure. That is the layered model we use in VL-e. It is certainly not complete and requires a lot more research.

COOK Report: And it is clear that you are thinking very creatively about how to follow existing ideas that have been used in software development to make operating systems more easily compatible with continuing improvements in microprocessor technology. Namely you could never afford to develop the science silos independently for each scientific discipline. You have to think collaboratively and to think about the processes that are needed in each discipline that can be generalized to a platform that is usable between disciplines. This is a high-level meaning of the process that you've been explaining. Furthermore you find that it also has added value within a discipline because it provides rationalization of parts of the research process itself as a form of pre standardization.

XI. Re-thinking American Infrastructure

What has to change

If you want to know how far the United States still has to go in terms of re-thinking, re-designing and actually building a national infrastructure for the 21st Century, a good place to start is http://www.infrastructurereportcard.org/. This website is run by the American Society of Civil Engineers to track the decline of America's crumbling infrastructure. Every year infrastructurereportcard.org publishes an annual grade for each of the critical infrastructure sectors that make up infrastructure for the United States. Then it averages those to assign an overall grade rating.

The bad news is that in 2009, the American Society of Civil Engineers gave the U.S. infrastructure a "D". Not quite failing, but almost. When it comes to Information Communications Technology infrastructure the news is even worse. There's no mention of ICT or the Internet at all. (But America did get a C+ in solid waste.)

Engineers still like to think about national infrastructure in terms of bricks, mortar, bridges and… solid waste.

And in the same way, when people think about a national ICT technology infrastructure they like to think about hardware, software, networks, storage, servers and databases.

19 years ago, when it was first published, the COOK Report focused on technology. First the bubble burst in 2000, and then the World Trade Centers came down on 9/11 in 2001. In both cases, failures related to technology were directly involved. There was more to technology than technology, so to speak. By 2003 the COOK Report covered Technology, Economics and Policy, and the disruption and chaos which seemed to ripple around that combination. It was clear the Internet and the growing fabric of services it carried were rapidly becoming the new infrastructure for the 21st Century. Unpacking that became the core activity for the COOK Report.

Most decision-makers prefer to think in terms of a world they can control, but the world we face is one of continually accelerating change in which multiple paradigms are changing at the same

time. Winston Churchill once described war as being a circumstance in which "all things are on the move simultaneously." What Moore's Law and Metcalfe's Law and maybe Renan's Law, too, all mean is that we no longer need war for the experience. Now all things are always on the move simultaneously.

The genius of the Dutch is that they chose to focus and invest in the one kind of national infrastructure which would assure them maximum opportunity and competitive advantage for the Netherlands, all who lived in it, learned in it and did business in it. In 1997 they embarked on the challenging task of building a national knowledge infrastructure supported by the most participatory and progressive ICT Research technology network, we think, in the world. The previously ten chapters have documented the extraordinary momentum they've created. Given we live in a world in which atoms are limited, but bits, evidently, are not, it was an excellent strategy, and many now view them as being the poster child for digital infrastructures.

The individual who seems to understand best the importance of technology for national development is Carlota Perez, and a favorite of the COOK Report. We devoted our entire May issue in 2008 to her work. She is both a visionary and a realist.

She titled a recent talk "The Future is Not Predictable Yet We Can Learn from History." This, of course, defines the challenge of implementing a successful infrastructure for the 21st Century. Much of what information technology does is find ways to make the past work with the future. And as change accelerates, so must agility. It takes everybody involved to optimize the process. The Netherlands has proceeded by means of an inclusive cooperative process.

The question is the same for every nation, big or small, poor or rich, developing or mature or in decline. How can we negotiate the transition from the past to the future? As a community? As a nation? And increasingly, as a species.

This is what Carlota Perez thinks about, talks about and writes about. Here is something she wrote in 2002:

"A golden age of worldwide expansion is possible. Making it happen will require thinking big, deciding wisely and acting boldly."

Carlotta Perez Epilogue: The World at the Turning Point Technological Revolutions and Financial Capital

It makes it seem as if this is a single decision, a single act, making the investment, etc. In other words, she describes — and governments and those who serve in them tend to think the same way —it as a matter of big dramatic steps.

But what we've learned from the Netherlands, in doing this special report and an earlier report last year, contradicts this. It is more than a matter of big steps. In fact, implementing a technology-based technology infrastructure is a very complex and long process that requires truly the best of intentions, but also considerable realism, pragmatism, patience and persistence. And we should not exclude the virtues of balancing consistency and flexibility ac-

complished at all stages and at every level.

Most importantly — and something that is rarely carried out by other countries — it must be negotiated with all stakeholders, again and again, over many years. People focus on the technical excellence of the new Dutch ecosystem. But the real genius is in the inclusiveness of the Dutch process, the carefulness of the governance, the attention to policy, execution… and total accountability.

So at this point, given success in the Netherlands, but a certain amount of confusion in the U.S., the logical question is why cannot the U.S. just use the models?

A History Lesson from the Tsar of Russia, Peter the Great

As Carlota Perez promised, history does offer guidance. Because this is not the first time the world has looked to the Netherlands for help in understanding technology and infrastructure. In the 1600s and 1700s, when national infrastructure was defined in terms of roads and canals, and no one in the world matched the Dutch for understanding canals, water and the first brilliant efforts at environmental engineering.

That is why, in 1697, Peter the Great, enormously talented Tsar of Russia, journeyed west to Amsterdam and Zaandam to learn Dutch technology. He studied with shipwrights and experts in many fields, to learn technology necessary for a national infrastructure, the building of ships and fortresses and how to design locks for canals.

He returned to Russia with as much technology knowledge as he could carry. This was a significant opportunity for Russia. But unfortunately he did not bring with the thing that made Dutch technology successful. He did not bring to Russia the Dutch political philosophy and the ability to work collaboratively. He focused on the technology, but not the practices that would allow it to foster innovation. Many believe this is why Russia never experienced an enlightenment.

And that is the lesson and the question for the United States.

The United States is at a turning point, it's progress faltering, its debt growing, its infrastructure beginning to crumble. Recall that in 2009 the American Society of Civil Engineers rated America's Infrastructure and gave it a "D".

So there is much the U.S. and the world have to learn from this new 21st Century knowledge infrastructure. But we have to make certain we pay attention to the whole and not simply the shiniest parts.

We need to begin with recognizing that the Dutch have a thoughtful, complete and balanced governance mode which takes into account the needs of the private sector as well as the public. And it is a very inclusive from another point of view.

Kees Neggers insists it is the users who run SURFnet. He likes to say: "all politics is local… technology organization works most effectively when it percolates bottom up from the plans of people who will pay for and use it."

But there is also thoughtful counsel, oversight and accountability at every stage, ending with the committee of "Wise Old Men" at the top. If you look at the chart on pages 21 they are the final step in advising before investment is made. And afterwards they are also the ones who examine if the promised results have been delivered.

The US Congress Office of Technology Assessment

Once upon a time in the U.S. we respected expertise and knowledgeable technology counsel: namely, The US Congress Office of Technology Assessment, known as the OTA.

When Newt Gingrich abolished the OTA in 1994 he abolished the US equivalent of the Committee of Wisemen. He opened the field to the K street Astroturf organizations that have made sure that US politicians remain ignorant of what goes on abroad. I worked full time from September 1 1990 through March 1 1992 on my own project at OTA. I can speak from first hand experience. OTA made sure that boundaries were crossed and the direction to congress were always here is what the technology does. Here are the implications and here is who "wins" and who "looses" if this legislation passes. In other words informed guidance in the public interest - laying out all the points of view. OTA was in existence for almost 30 years and was a model for many European approaches.,

No more. Its demise opened the door for the incumbents to use their huge cash flow in their very narrow economic interest. With no John Kennedy asking "not what your country can do for you but what you can do for your country' and only worship of

private good and free market for the past 30 years. With the cold war "won" we have again lost the ability to think large. A national agenda comparable to sending men to the moon is not attainable in the midst of partisan rancor.

Examining the Public Good versus Private Opportunity

While public versus private good is a useful metaphor to think of in terms of defining telecommunications access, because telecommunication access costs something to provide, it may not fit very snugly into the category of narrowly defined public good. See: http://en.wikipedia.org/wiki/Public_good

Nevertheless, roads, water-works and electricity also are not free and society does provide those on a subsidized basis and one not motivated by pure profit. The logical and absolutely necessary answer to the "fit" is infrastructure and in the USA the 1934 act recognized voice phone service as infrastructure and acted accordingly.

However on the way to Googin's paradox, the picture of what was and what was not infrastructure became muddied. Because the telephone companies became the

conduit for internet to the masses during the 1990s during a period marked globally by deregulation and private good free market mania - the emergence internet technology, although potentially a general purpose technology, became locked within a mesh that is seen -- in the USA and to a lesser extent in the rest of the world -- as the realm of private good financial capital and therefore something in which governments should not be involved.

Because of the predominance of private good thinking we keep more to our narrow silos and are very much less equipped to take advantage of the thinking on which the Netherlands builds.

And therefore extending Googin's paradox to the whole world -- private good thinking keeps you not only from sharing the benefit of optical technologies…. it also locks you into not invented here way of looking at reality and handicaps your ability to collaborate.

Changes at the FCC

The jury is still out at the FCC. But there is reason to be hopeful. There is an international group at the FCC assigned to work on a comparison of twenty foreign countries' broadband policies from

a community network and regulatory point of view. Involved are Narda Jones, Chief Strategic Analysis and Negotiations Division, International Bureau and Irene Wu Yahoo! Fellow in Residence Director of Research, SAND-MNIA International Bureau Federal Communications Commission.

I have personally talked with some of the National Broadband Plan team. I have been extremely encouraged by the experience. The political process didn't get all of them on board until roughly half the assigned year was gone. Another example of the situation where the U.S. has built up a destructive momentum, which is not proving so easy to reverse.

What has to change?

The ICT ecosystem and culture of the Netherlands is, of course, reflected in many technology communities in the U.S. We do not have a great technology commons here, but we have literally thousands of smaller commons blooming everywhere in the United States.

To create a successful national infrastructure for America for the 21st Century will require a fundamental re-examination of American values and governmental. So far, it appears this is a very

large obstacle. But this does not mean that all of Netherland's innovations will be ignored.

The E-Science Suite Is Not Just for Science

Initiatives like GLIF are, of course, already been used for international collaboration for doing science. The scientific community, will as it has always done, continue to have a more open-mind than most of American Enterprise. But what will be adopted, the COOK Report believes, is the new Netherland's impressive new generation of resources, tools and process for e-Science. Internally, we've been calling it the e-Science Suite. It's our name, not one the team in the Netherlands uses. We fell into using it because their approach was to, step-by-step, create a growing suite of generic middleware, shared tools, support for collaborative solution solving, and various modalities to support a more multidisciplinary approach. What Bob Hertzberger and the VL-e team have done in Amsterdam is what we call in America "heavy lifting".

The problems the e-Science Suite solves, in terms of making it easier and more affordable for new domains of science and communities to begin work with and share enormous databases, are not

trivial. It will be widely discussed and often influential. Many will chose to join these new collaborations. As a result they will be able to take advantage of a more sophisticated, more shareable infrastructure. This will allow them to access, share and work with enormous data streams and structures.

We do not believe the e-Science Suite tools and processes will be used only for science. It is well suited, as development continues, for many 21st Century style activities in which competitive advantage and success depend on a deepening and more seamless collaboration.

Nobody is calling VL-e's e-Science Suit the 5th Paradigm. But over time, we wouldn't rule it out.

The Netherlands vs the USA in an education that is part of the larger knowledge infrastructure

An example of the disparity is ICT support for education in the context of overall knowledge infrastructure.

The Netherlands have designed something like a Linux OS for all aspects of the Dutch education, research, technology and product development. It is a mechanism designed to fit the intel-

lectual and subject matter strengths of the Netherlands. It is designed to create a generic commons overlaid on a network fabric with the basic tools that individual disciplines need to ply their trade. It is being done likely for the first time through a communications web of planning that gathers the stake holders together with the intent of tuning them into collaborators rather than competitors.

Is an ICT research Infrastructure for a nation tantamount to the creation of a knowledge framework for a society to experiment in new forms of collaboration, problems solving invention, creativity. The Netherlands has built a multi dimensional matrix for the application at one level of Renan's law of connecting everything and Baken's metaphysics of holon's. Building as infrastructure means that the Netherlands organizes what it does in new ways and does so it more efficiently. Approaching it as infrastructure helps pry participants out of disciplinary silos.

In the USA what we have is too often pork and plunder when it comes to ICT support for education. (The following url takes you to the archtypical story of IBM supplying Internet for El Paso Schools, rigging bids, building too complex a system, charging too much, and after the bank

was broken, pulling out the infrastructure and leaving the schools without connectivity.) www.nytimes.com/2004/06/17/politics/17computer.html?pagewanted=all

Unfortunately this is only one of the problems blocking progress. Read on.

The power of Research Networks in the Netherlands.

It has lead to a not fully understood dichotomy where research networks are transforming what they touch in a way that the current public stasis renders impossible for the rest of us to grasp. Not connected to them we cannot experience them first hand. We in effect don't know what we are missing. In the past, the catch-word has been patience, it will come to you. Well this time the truth is it won't. The incumbent business model forbids it.

The question is how to break out?

Because if we don't, we are deprived of not only of the SURF sponsored technology transfer described earlier in this paper : http://www.jisc.ac.uk/whatwedo/campaigns/res3.aspx and The Future of higher Education: http://bit.ly/5ejtEl.

Moreover to ensure 100% return on investment for the incumbents when ordinary savings accounts may pay one half of one per cent is not the idea of equal justice for all on which the Unites States was founded. Secondary schools will also remain locked out of the benefits that have been described in this paper. Why because the FCC created the school and libraries fund in such a way that it operates for the benefit of the incumbents scarcity model. They are forced to become customers for the incumbents and prevented from using universal services funds to establish their own service at less cost. They were forbidden to offer their own services setting up a two billion dollar a year feeding trough for the incumbents. The consequences were precisely what Dave Hughes the "cursor cowboy'" predicted when the program was established more than a dozen years ago. The El Paso story above and thousands more were inevitable.

If R and E networks succeed in bringing this future to some campuses but not to all education, I contend it is an invitation to another separate and unequal case analogous in the United States to the Supreme Court 1954 decision abolishing school segregation. Don't misunderstand me. The R and E networks must succeed and the incum-

bent must accept structural separation or have it imposed by the government in the name of a sustainable future for society. Just as schools were ordered desegregated in 1954.

And the pace of change increases. It never slows. Just when things look darkest come signs of the dawn. In 2005 David Isenberg held the first of five conferences on the political issues around connectivity and access to resources in Washington DC. The theme was Freedom 2 Connect.

Five years later, on Thursday January 21 2010 Hillary Clinton enshrined these three powerful words into official US policy.

The Freedom to Connect

"The final freedom I want to address today flows from the four I've already mentioned: the freedom to connect – the idea that governments should not prevent people from connecting to the internet, to websites, or to each other. The freedom to connect is like the freedom of assembly in cyber space. It allows individuals to get online, come together, and hopefully cooperate in the name of progress….

The spread of information networks is forming a new nervous system for our planet. When something happens in Haiti or Hunan, the rest of us learn about it in real time – from real people. And we can respond in real time as well…

As we sit here, any of you – or maybe more likely, any of our children – can take out the tools that many carry every day and transmit this discussion to billions across the world.

Let us make these technologies a force for real progress the world over. And let us go forward together to champion these freedoms."
http://blog.austinheap.com/transcript-of-sec-clintons-speech-on-internet-freedom/

Two days later writing from the Netherlands Jaap van Till elaborated on what this means.

"The young of this world, emancipated by online networking, do no longer want old men to tell them what to think."

"What we must help teach them is how to think and how to collaborate via cell phones, tablets and laptops. Strong community network links and 'weak links' will appear between them and stabilize the world. No dictatorial government, bureaucrat, spin doctor, priest or mullah will be able to stop the kids doing that."

Having gotten "freedom to connect" accepted as part of the bed rock foundation of official US policy is a very important achievement on the part of David Isenberg. But the operative word is connect and that leads to Renan's law. Everything wants to be connected. Everything works better when connected and that the absence of connection is death.

But We End with Hope

Renan's law however is becoming more and more evident. Augmented reality and something called "the sixth sense" are rapidly emerging technologies. As a result our movements in the physical world will blend with the digital world with much less effort and very quickly the computer effectively vanishes. Computing in effect will merge with the physical world. Newspapers for example will show live information you will be able to pick up a book in a book store and your interface will show you and tell you whatever you would like to know about the book. In effect "these devices by making the machines disappear wind up helping us to be more human so that we do not wind up being machines sitting in front of other machines."
http://www.ted.com/talks/view/id/685

Pranav Mistry leads a group at the media lab and India, that is developing this technology.   The hardware is commodity.   The code is open source. Because they cannot be locked in silos, the walls of scarcity we have been discussing in this book will either dissolve or become irrelevant.   Renan's law will triumph because things indeed work better when connected. The Netherlands has built the receptive environment for what Renan and Mistry describe. The United States must catch up.

Contents

COOK REPORT
ON INTERNET PROTOCOL
FEBRUARY 2009

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?

"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 ex

amine 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 De-Fanti, 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 is-

sue. 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 telecommuni-

cations 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,

Amsterdam Science Park Watergraafsmeer   GigaPort

SURFnet

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 connec-

tions 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 interworking 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 be-**

gin 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 de-

scribe 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 prob-

lem. 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.

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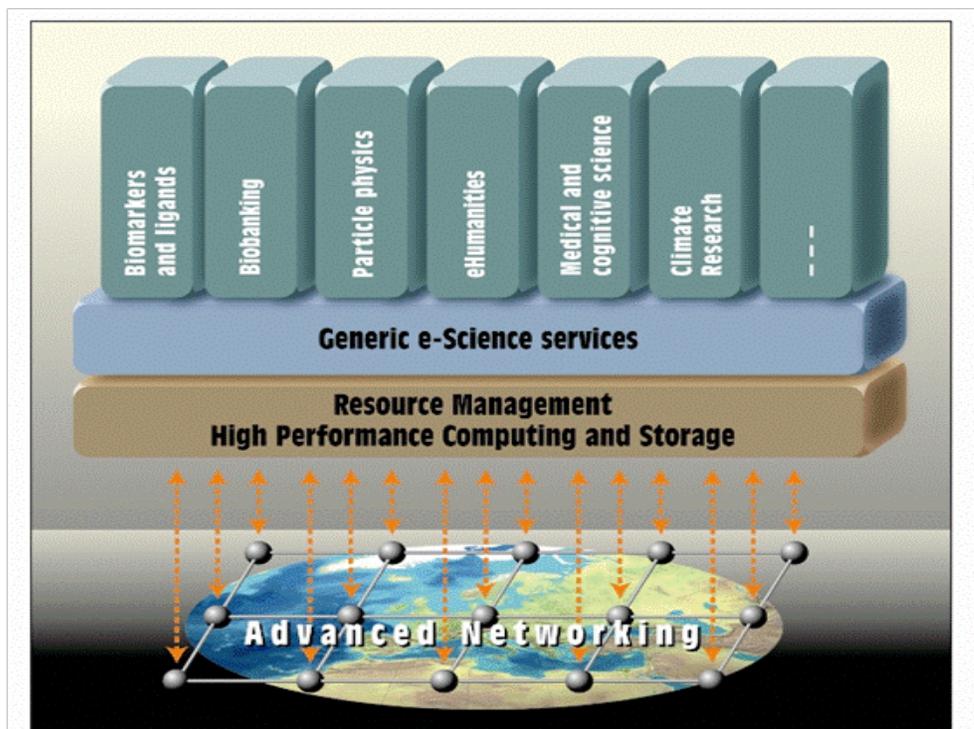
Virtual laboratory for e-Science, *Bob Hertzberger, Henri Bal* et. al.

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&pag

ewanted=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
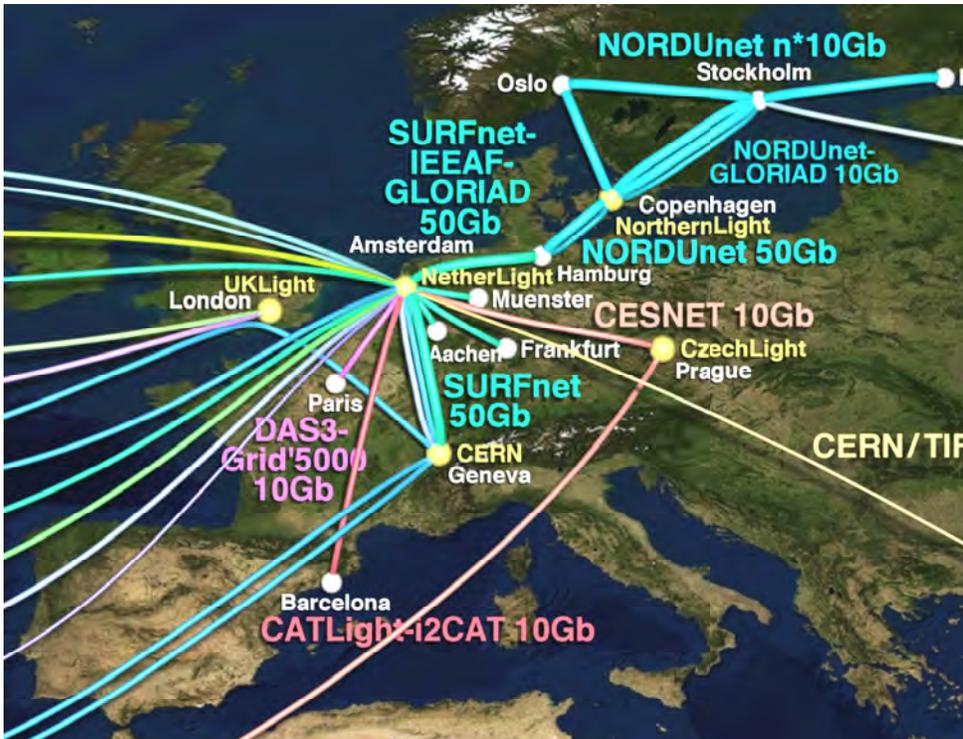
applications to meet world-wide.

COOK Report: While in theory then an American innovation 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?

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. NOR-DUnet and NLR are also taking part in this procurement and will be getting their own

Kees Neggers at Sc08 in Austin

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

System &Network Engineering @ UvA
update 2008

- group has 4 sections
  - Advanced Networking (GP, EU, TNO)
    - Paola Grosso
  - Security (GP, EU, VL-e, SurfWorks)
    - Guido van 't Noordende
  - Sensor Grids - Intelligent networks (TNO)
    - Rob Meijer
  - Master SNE education (GP)
    - Karst Koymans
- 25 people - 19 fte
- Home @ Science Park Amsterdam, co-located with:
  - NIKHEF (together with SARA LHC Tier-1 center, BigGrid)
  - SARA (SN6-NOC, NetherLight, SN6-core location, LightHouse)
  - AMS-IX
  - UvA Science faculty (Dutch e-Science program VL-e)

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

GigaPort - Plans 2004-2008

1. **Hybrid networking structure**
   - Network Architecture
   - Optical Internet Exchange Architecture
   - Network Modeling <NDL, Pathfinding>
   - Fault Isolation
2. **Network transport protocols**
   - UDP - TCP
   - Protocol testbed
   - LinkLocal Addressing
3. **Optical networking applications**
   - StarPlane
   - eVLBI
   - Smallest University for proof of concepts
   - CineGrid
   - CosmoGrid
4. **Authorization, Authentication and Accounting in Networking and Grids**
   - AAA & schedule server
   - WS security
   - Multi domain token based implementations
   - Cross domain LightPath setup
5. **Testbed LightHouse, SC0X, iGrid, GLIF, OGF, Terena, ...**

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 use-

ful 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 ca-

pacity. 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 us-

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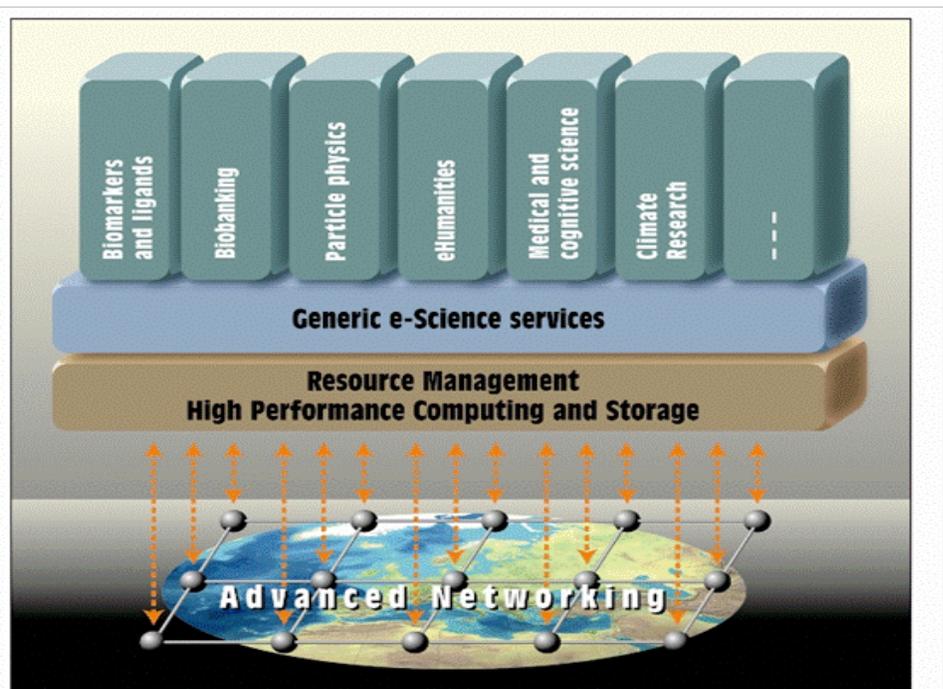
Virtual laboratory for e-Science, *Bob Hertzberger, Henri Bal* et. al.

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.
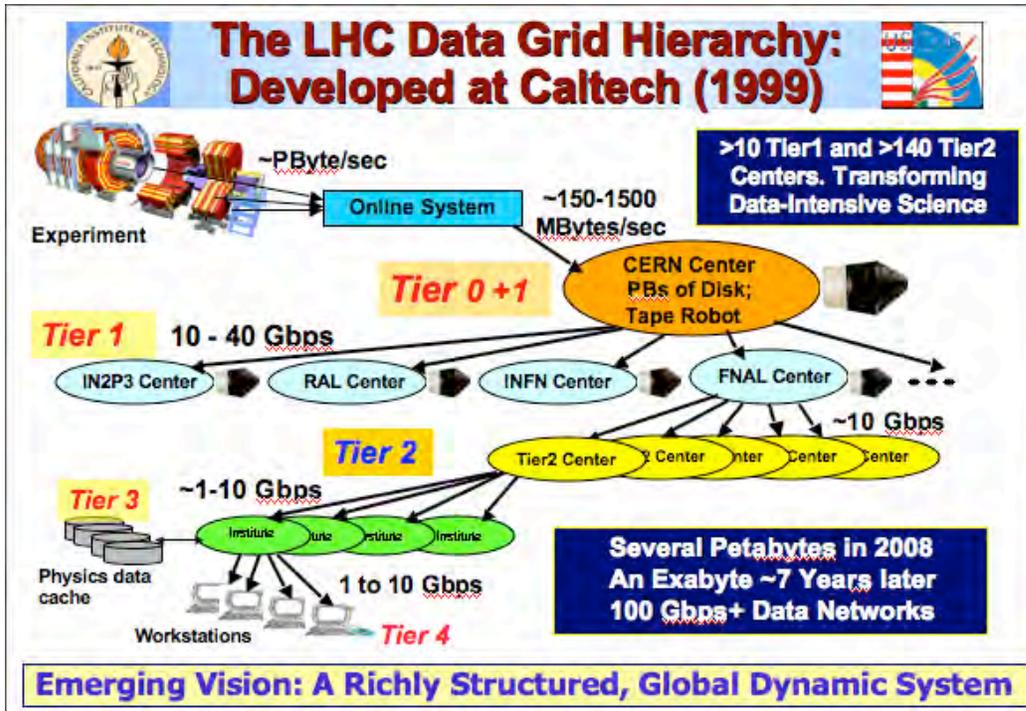
and/or send the data to tier-2 locations and local universities that will process and interpret the data and use the results to do science. Results are then send back to the tier-1's for archival and access by scientists purposes. [**Editor**: The slide shown is from Harvey Newman's September 2008 "deck."]

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.

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
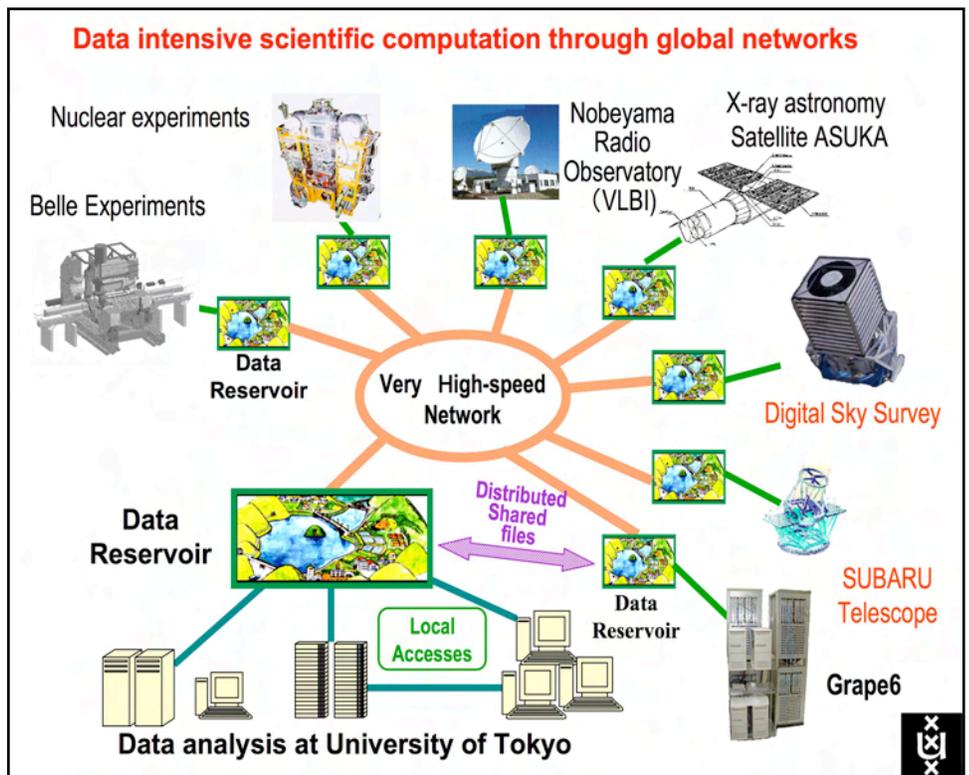
WiVR: a Window into Virtual Reality

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 posi-

tion 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.

CineGrid@SARA
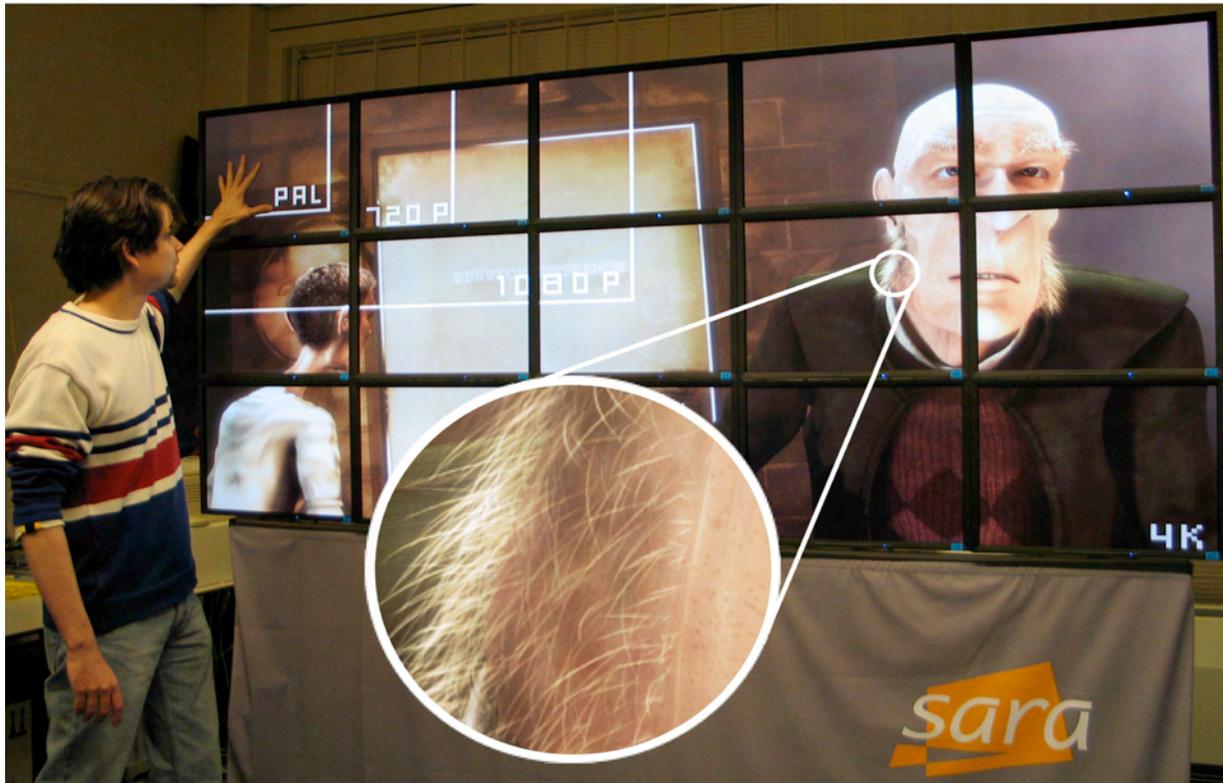
The slide above shows the visualization of cinematic SHD content (3840 by 2160 pixels) on a tiled display setup at the supercomputing center in Amsterdam.
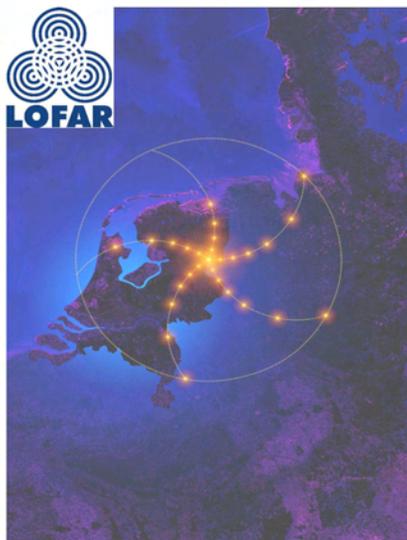
LOFAR: the Very Long Baseline Interferometer radiotelescope array.

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.

Towards Hybrid Networking!
- Costs of photonic equipment 10% of switching 10 % of full routing
  – for same throughput!
  – Photonic vs Optical (optical used for SONET, etc, 10-50 k$/port)
  – DWDM lasers for long reach expensive, 10-50 k$
- Bottom line: look for a hybrid architecture which serves all classes in a cost effective way
  – map A -> L3 , B -> L2 , C -> L1 and L2
- Give each packet in the network the service it needs, but no more !

L1 ≈ 2-3 k$/port

L2 ≈ 5-8 k$/port

L3 ≈ 75+ k$/port

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: 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.

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

Trends
- We have made baby-steps on the path to optical networking
  – Still many mails and phone calls
- See several trends:
  – lambda's get fatter and cheaper
  – photonic technology cheap per bandwidth
  – embedded computation capacity increasing
  – latency and high bandwidth congestion avoidance conflict
  – ethernet is getting circuit properties (PBT)
  – applications need more and more predictable behavior

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.

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
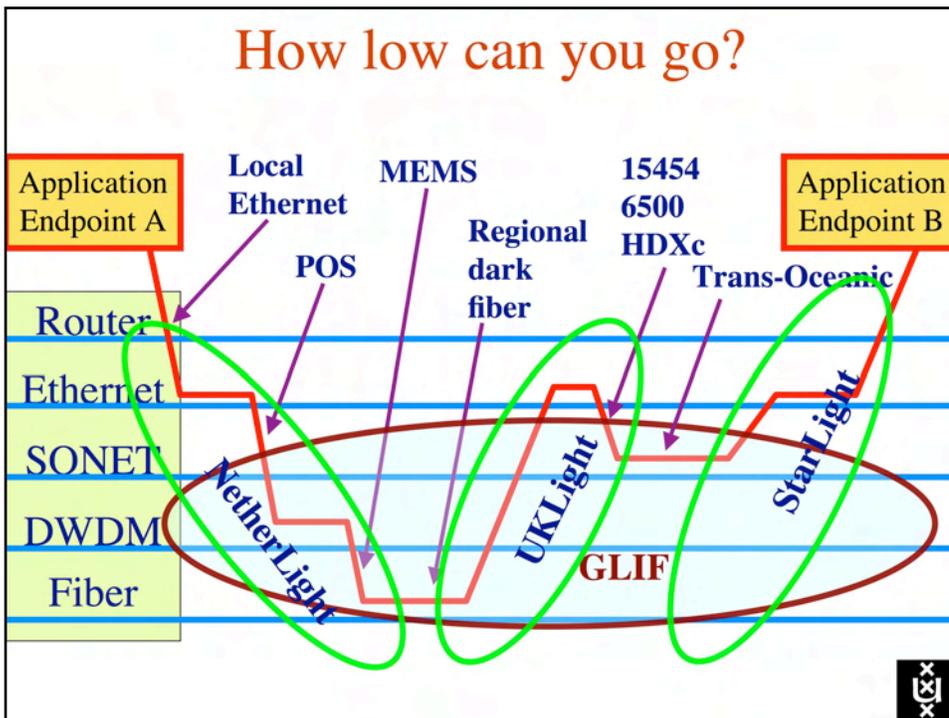
Your Network Application Must Use the Appropriate Transport Layer

GLIF 2008 — Visualization courtesy of Bob Patterson, NCSA. Data collection by Maxine Brown.
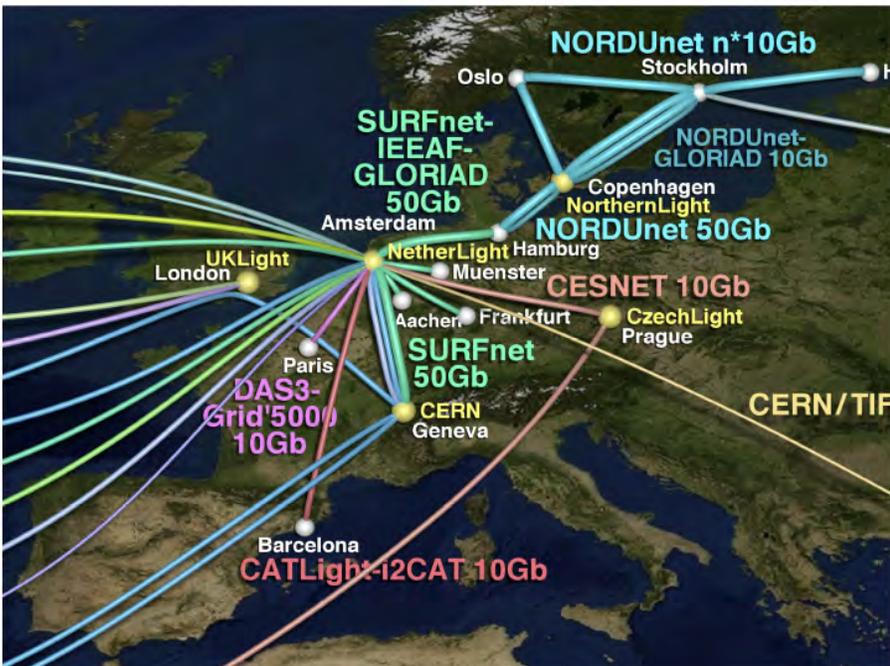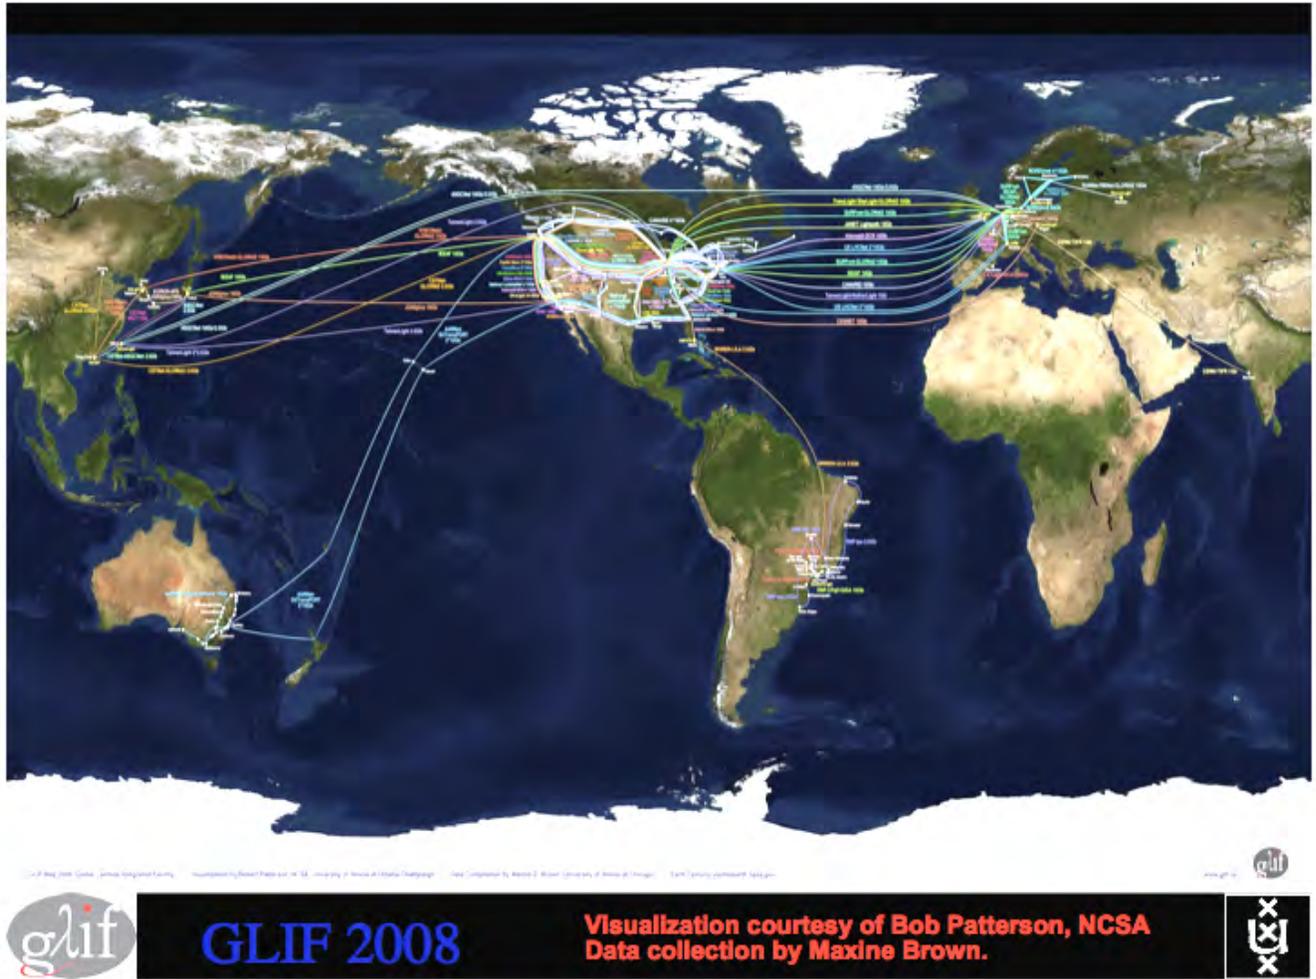
the Netherlands, NetherLight and Amsterdam. First and to the left you can see the connectivity and, in the second graph at the top of the next page and overlaid on the connectivity map, the activities that flourish around NetherLight in Amsterdam.

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 infra-structure: Visualization, Data Storage and manipulation, grid and cloud computing and supercomputing. The map below shows the SURFnet fiber infrastructure.

In The Netherlands SURFnet connects between 180:
- universities;
- academic hospitals;
- most polytechnics;
- research centers.
with an indirect ~750K user base

Red crosses = StarPlane

~ 6000 km scale comparable to railway system

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

De Laat: perhaps about four or five
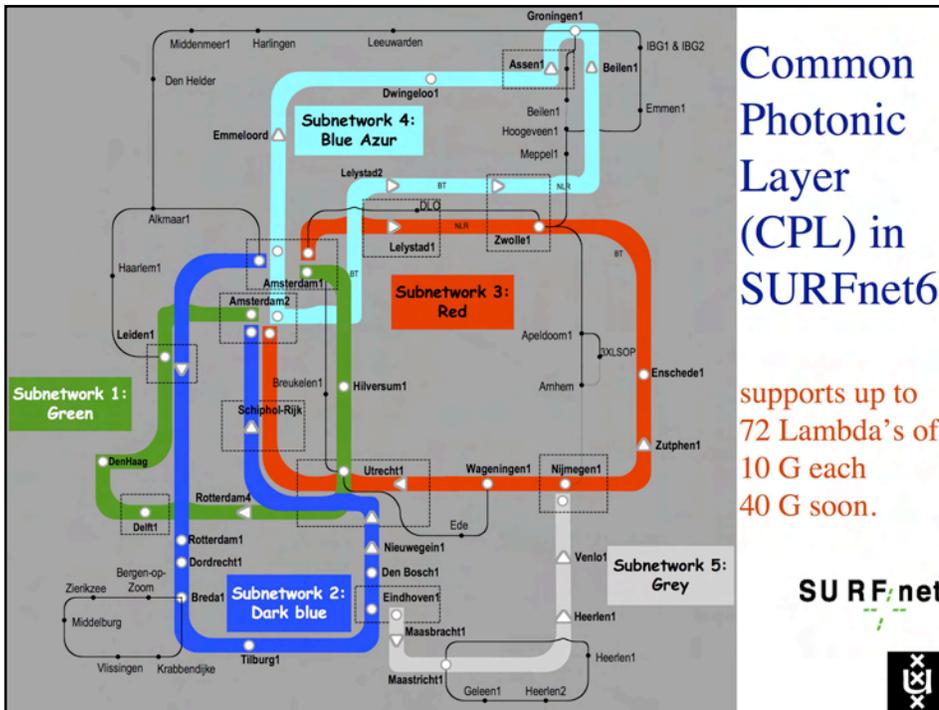
resources and to create multiple overlay networks, each with a different logical topology. The novelty of this project is that it does give this flexibility directly to the applications by allowing them to choose the logical interconnection topology in real time, ultimately with subsecond switching times. 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/

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 Star-Plane 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

We do this experiment with five distinct clusters at 4 different Universities - hence
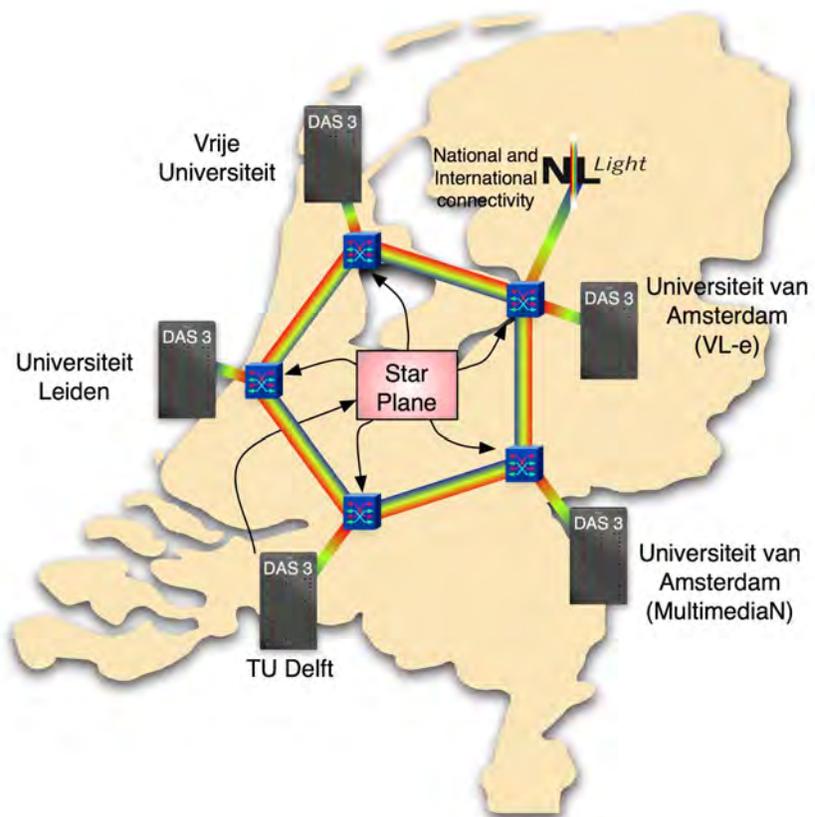
QOS in a non destructive way!

- **Destructive QOS:**
  - have a link or λ
  - set part of it aside for a lucky few under higher priority
  - rest gets less service

- **Constructive QOS:**
  - have a λ
  - add other λ's as needed on separate colors
  - move the lucky ones over there
  - rest gets also a bit happier!

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.

Module Operation

input and output fibres

common   port 1   collimating lens

diffraction grating

> this schematic shows
  - several input fibres and one output fibre
  - light is focused and diffracted such that each channel lands on a different MEMS mirror
  - the MEMS mirror is electronically controlled to tilt the reflecting surface
  - the angle of tilt directs the light to the correct port

> in this example:
  - channel 1 is coming in on port 1 (shown in red)
  - when it hits the MEMS mirror the mirror is tilted to direct this channel from port 1 to the common
  - only port 1 satisfies this angle, therefore all other ports are blocked

λ_1 λ_n

MEMS mirror array
(1 pixel per channel)

ref Eric Bernier, NORTEL

COOK Report: Is this light-path 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 to-

pologies. 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 light-path 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-distorted signal gets distorted through transportation it is normal again. You

Dispersion compensating modem: eDCO from NORTEL
(Try to Google eDCO :-)

sender → receiver

$h(t)$

$T(f)$
transport as function of frequency T(f)

$F^{-1}(F(h(t)).T(f))$

Solution in 5 easy steps for dummy's :
1. try to figure out $T(f)$ by trial and error
2. invert $T(f) \to T^{-1}(f)$
3. computationally multiply $T^{-1}(f)$ with Fourier transform of bit pattern to send
4. inverse Fourier transform the result from frequency to time space
5. modulate laser with resulting $h'(t) = F^{-1}(F(h(t)).T^{-1}(f))$

sender with cpu, mod → receiver

$F^{-1}(F(h(t))*T^{-1}(f))$

$T(f)$

$F^{-1}(F(F^{-1}(F(h(t)).T^{-1}(f))).T(f)) \to h(t)$

(ps. due to power ~ square E the signal to send **looks** like uncompensated received but is not)

GRID Co-scheduling problem space

CPU   DATA   Lambda's

Extensively under research

New!

The StarPlane vision is to give flexibility directly to the applications by allowing them to choose the logical topology in real time, ultimately with sub-second lambda switching times on part of the SURFnet6 infrastructure.
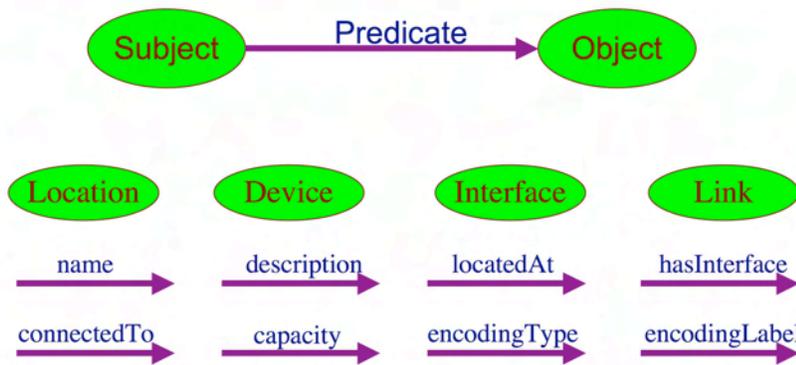
Austin (TX). The example is from SC06.

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.

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

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 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

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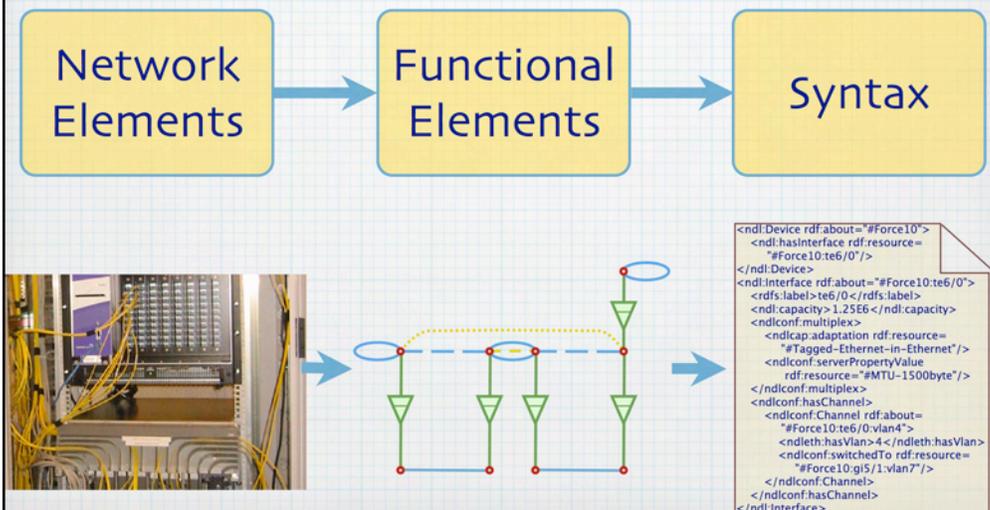
COOK Report: And conceptually this is down in the engine room of the path finding process for which the light table on the show floor represents the graphical user interface?

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?

NetherLight in RDF

```
<?xml version="1.0" encoding="UTF-8"?>
<rdf:RDF xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
     xmlns:ndl="http://www.science.uva.nl/research/air/ndl#">
<!-- Description of Netherlight -->
<ndl:Location rdf:about="#Netherlight">
     <ndl:name>Netherlight Optical Exchange</ndl:name>
</ndl:Location>
<!-- TDM3.amsterdam1.netherlight.net -->
<ndl:Device rdf:about="#tdm3.amsterdam1.netherlight.net">
     <ndl:name>tdm3.amsterdam1.netherlight.net</ndl:name>
     <ndl:locatedAt rdf:resource="#amsterdam1.netherlight.net"/>
     <ndl:hasInterface rdf:resource="#tdm3.amsterdam1.netherlight.net:501/1"/>
     <ndl:hasInterface rdf:resource="#tdm3.amsterdam1.netherlight.net:501/3"/>
     <ndl:hasInterface rdf:resource="#tdm3.amsterdam1.netherlight.net:501/4"/>
     <ndl:hasInterface rdf:resource="#tdm3.amsterdam1.netherlight.net:503/1"/>
     <ndl:hasInterface rdf:
     <ndl:hasInterface rdf:
     <ndl:hasInterface rdf:
     <ndl:hasInterface rdf:
     <ndl:hasInterface rdf:
     <ndl:hasInterface rdf:
     <ndl:hasInterface rdf:
     <ndl:hasInterface rdf:
     <ndl:hasInterface rdf:
```

```
<!-- all the interfaces of TDM3.amsterdam1.netherlight.net -->

<ndl:Interface rdf:about="#tdm3.amsterdam1.netherlight.net:501/1">
          <ndl:name>tdm3.amsterdam1.netherlight.net:POS501/1</ndl:name>
          <ndl:connectedTo rdf:resource="#tdm4.amsterdam1.netherlight.net:5/1"/>
</ndl:Interface>
<ndl:Interface rdf:about="#tdm3.amsterdam1.netherlight.net:501/2">
          <ndl:name>tdm3.amsterdam1.netherlight.net:POS501/2</ndl:name>
          <ndl:connectedTo rdf:resource="#tdm1.amsterdam1.netherlight.net:12/1"/>
</ndl:Interface>
```

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

NDL Generator and Validator

see http://trafficlight.uva.netherlight.nl/NDL-demo/

Multi-layer extensions to NDL
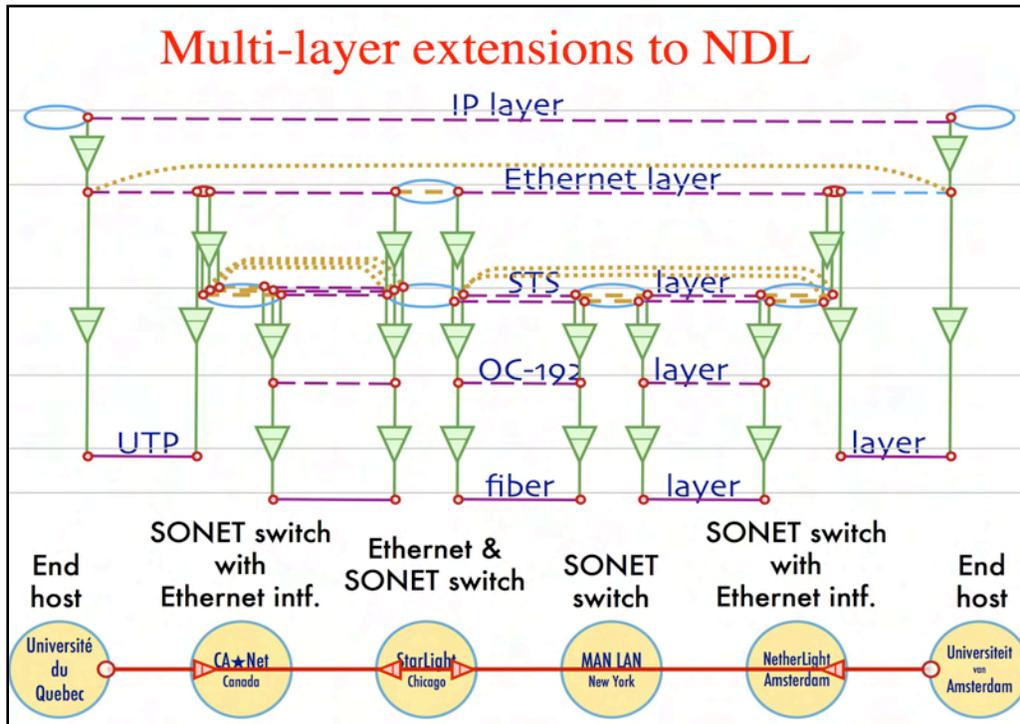
a visualization of SURFnet6 based on a NDL description. Although the visualization has lots of room for improvement, it gives confidence that the model works.

Network Description Language as Routing Pro-


NDL SN6 Visualisation
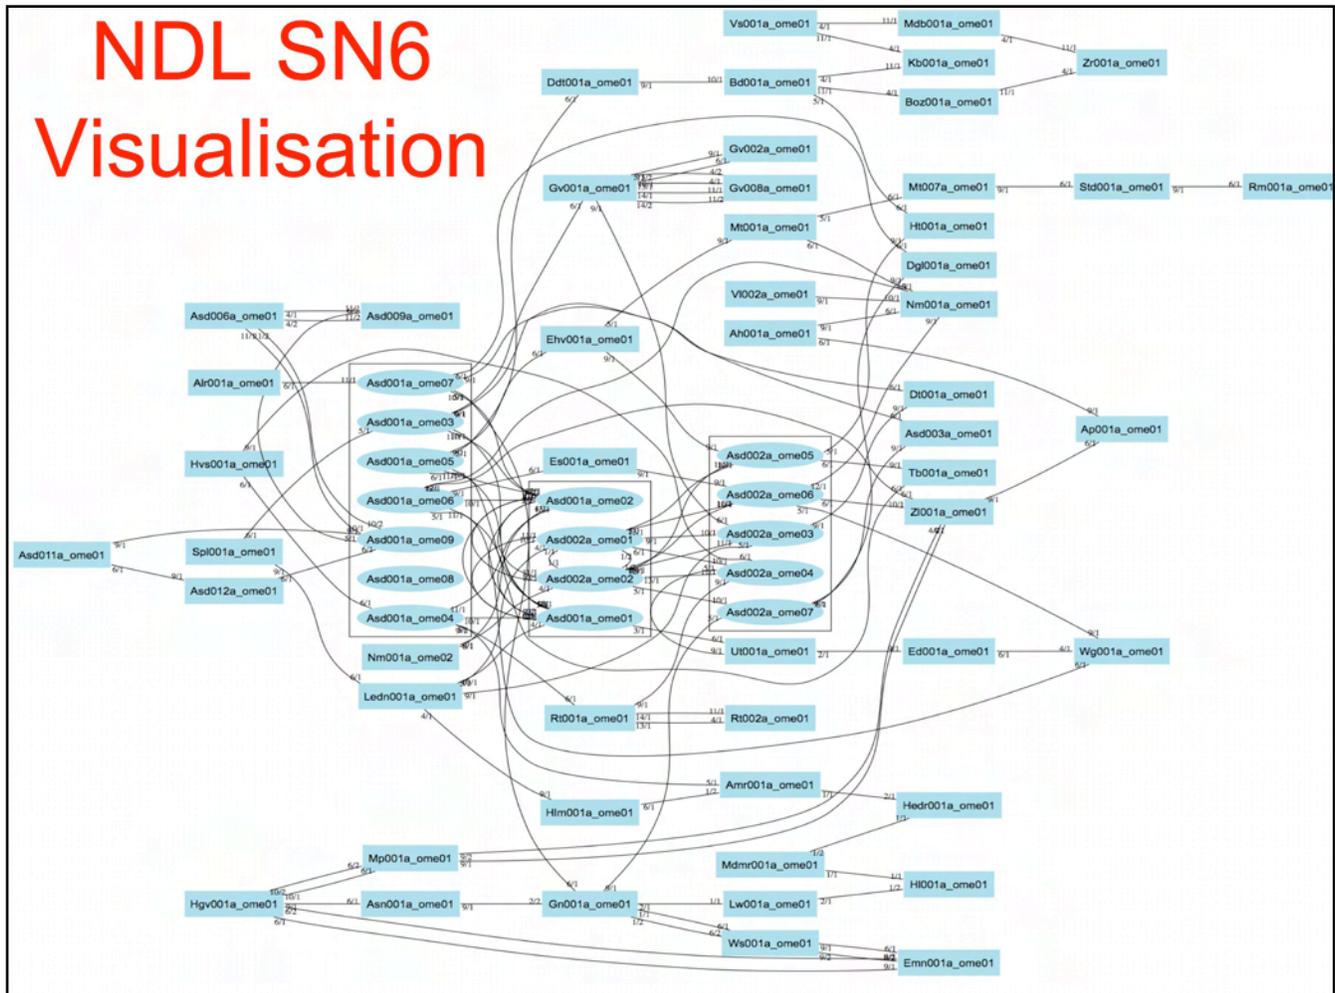
tocol

In the next two slides we demonstrate the multi domain multi layer routing power of the model in a

quires a loop to allocate a gigabit lightpath from source to destination.

COOK Report: So in the first slide below you can not go

mock-up network that re-

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 Star-Light 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 MAN-LAN 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.
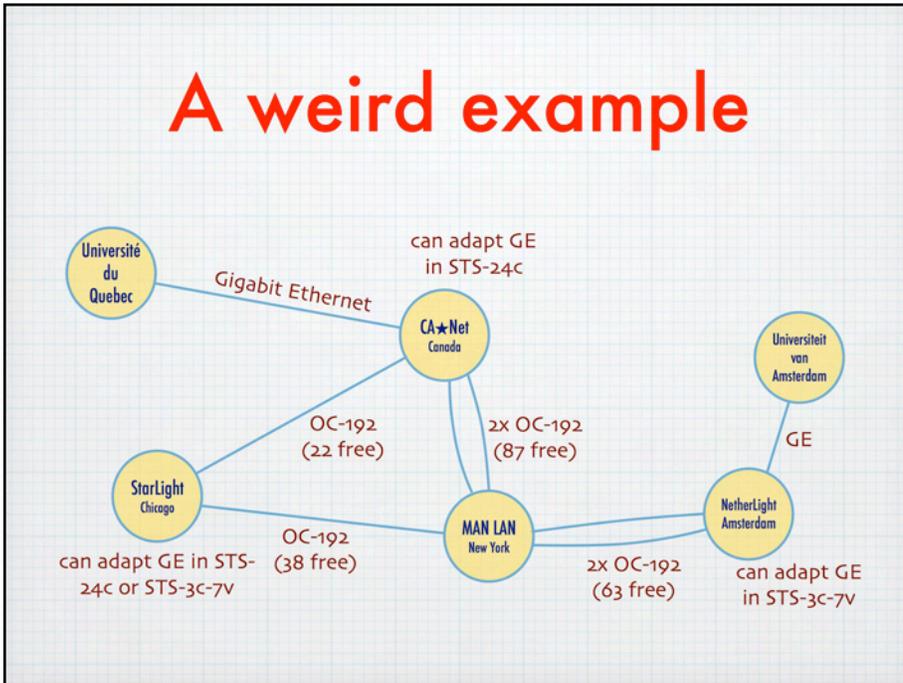
COOK Report: This is like network state information?

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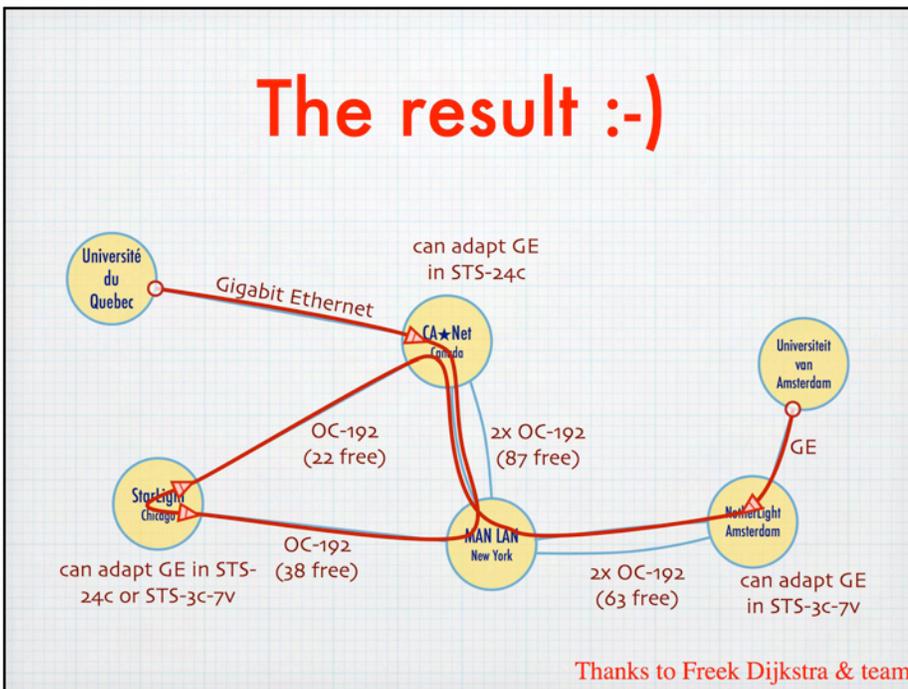
And on the top of the next page is a demonstration where the protocol goes through the steps of the state machine.

By tracing the course of the Green line from the left and

MultiDomain MultiLayer pathfinding in action
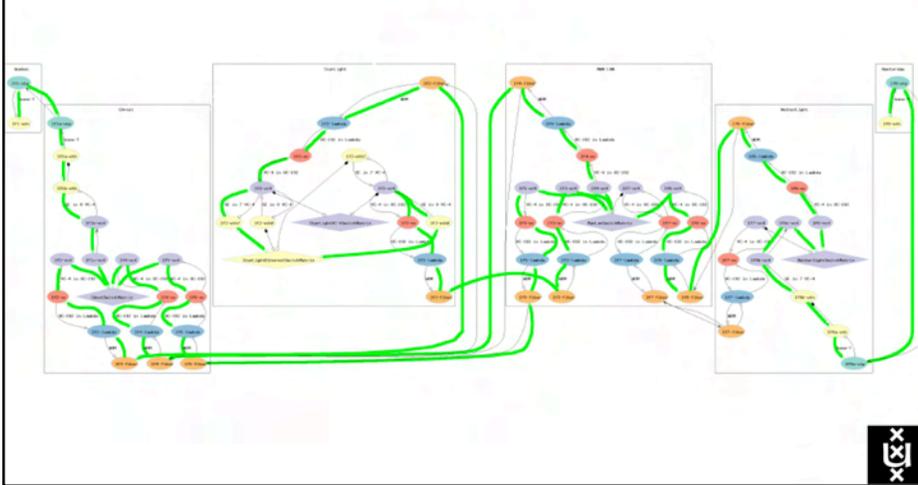
the right you can see how the network protocol tries all possibilities and finds its way through the maze. We are submitting NDL to standardization organizations.

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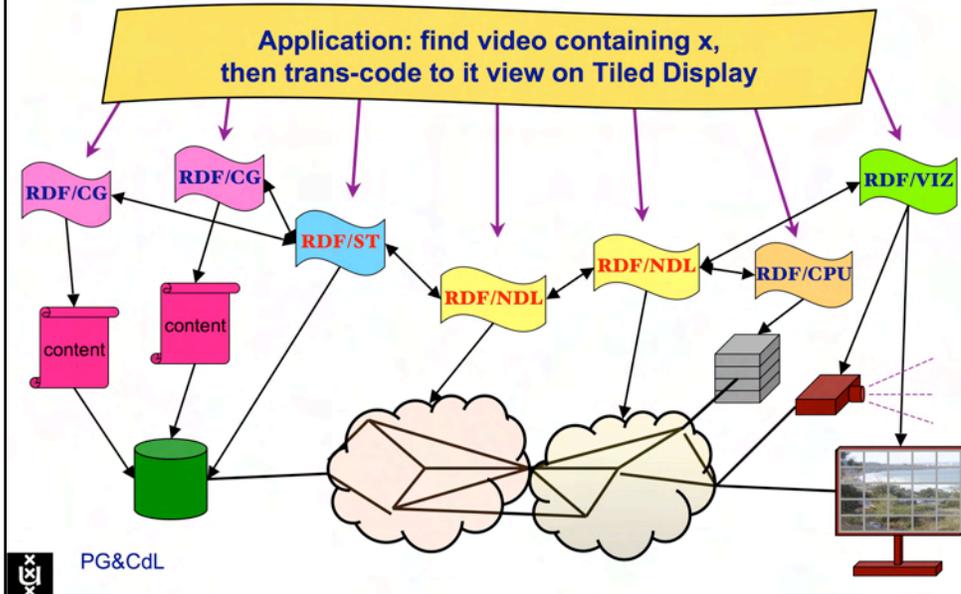
COOK Report: Who else in the world is doing anything like this?

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."

RDF describing Infrastructure

Application: find video containing x, then trans-code to it view on Tiled Display

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 sec-

ond! 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?
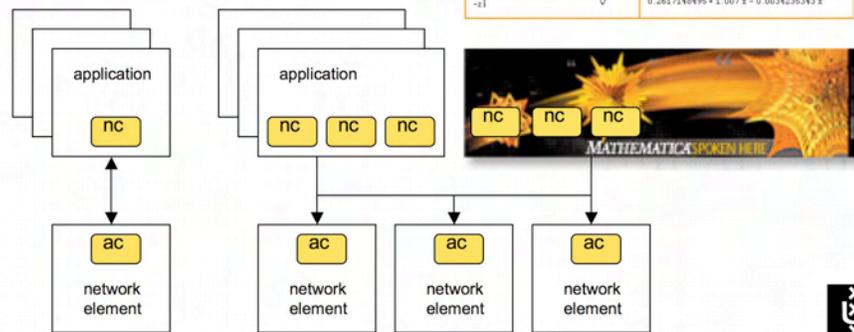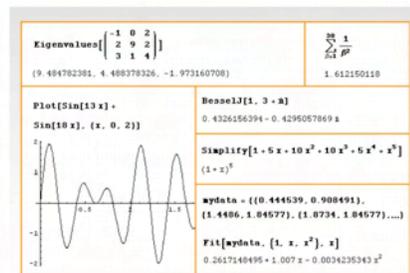
De Laat: Correct. I consider it to be a Tbit/s network if all that capacity can be used by one application for its purposes. For that one needs programmable networks. So we asked the question of what constitutes a terabit network in a different way. Think back to teraflop computing. People started out 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 func-

Mathematica enables advanced graph queries, visualizations and real-time network manipulations on UPVNs

Topology matters can be dealt with algorithmically
Results can be persisted using a transaction service built in UPVN

ref: Robert J. Meijer, Rudolf J. Strijkers, Leon Gommans, Cees de Laat, User Programmable Virtualiized Networks, accepted for publication to the IEEE e-Science 2006 conference Amsterdam.

Functional building blocks

Control Interface (protocols API's)

Use Interface

Application

Application layers

Network layers

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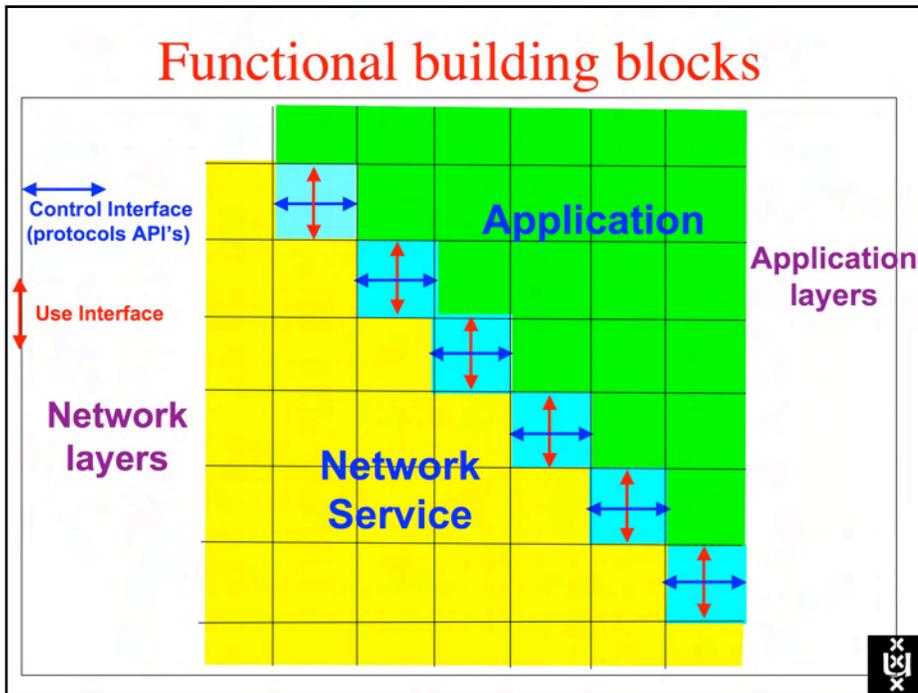
Finally, in dealing with ICT infrastructure, power consumption is a big issue. The more switching you can do in the photonic domain, the more power you can save.

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

tions 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.

Power is a big issue

- UvA cluster uses (max) 30 kWh
- 1 kWh ~ 0.1 €
- per year   -> 26 k€/y
- add cooling 50%   -> 39 k€/y
- Emergency power system   -> 50 k€/y
- per rack 10 kWh is now normal
- YOU BURN ABOUT HALF THE CLUSTER OVER ITS LIFETIME!

- Terminating a 10 Gb/s wave costs about 200 W
- Entire loaded fiber -> 16 kW
- Wavelength Selective Switch : few W!

Cees de Laat at SC08 in Austin

each of those interfaces consumes 200 Watts.

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.

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.

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

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 manipu-

lated 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 con-

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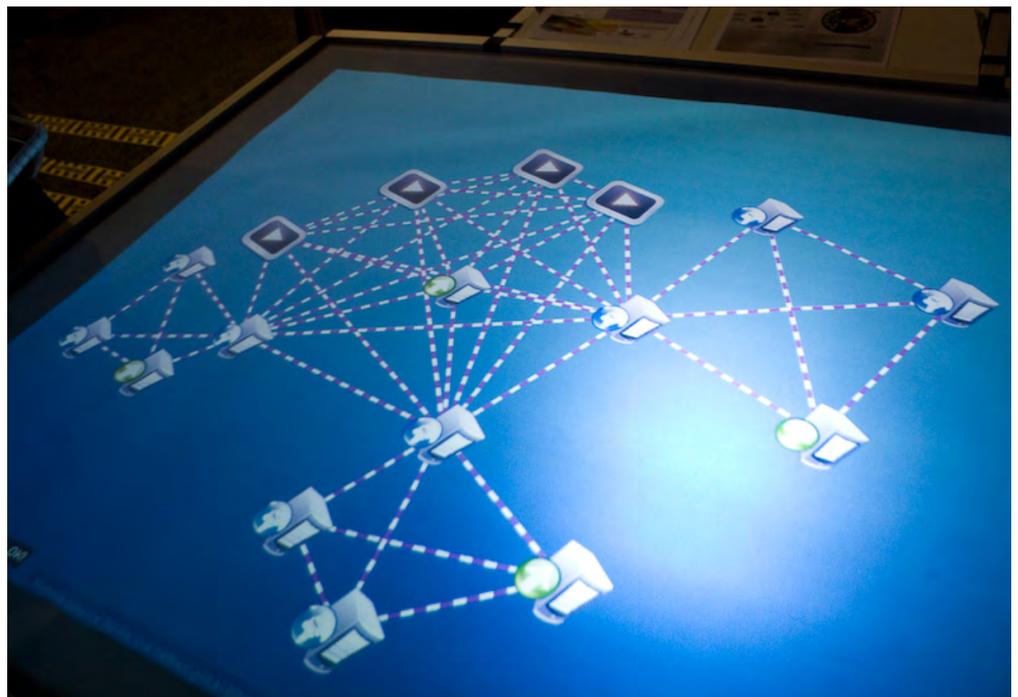
Figure 2. Network visualization after discovery. The dashed lines, when monitoring is enabled will indicate bandwidth (line width), delay (percentage of dashing) and jitter (randomize in dashing).

nipulations as the touch ta-ble. We will come back to this aspect later.

COOK Report: What do these icons represent?

Strijkers: The icons repre-sent 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 visual-ized 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 pre-view 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 Pro-ject and the son of Cees de Laat made one of the Cine-Grid movies. All the movies stream continuously stream in high definition, but to un-known 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 measure-ments. 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 be-comes larger than 1, the nodes will light up red, alert-ing 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 be-low.]

Strijkers: Yes. If we make a path, what will happen, we send a request to the control-ler asking to create the path that we just dragged with our finger. The controller will send the request to a com-piler to generate the com-mands and forwarding ex-pressions for provisioning the nodes. The results will be passed on to a transaction monitoring, which executes a distributed transaction to load the commands and ex-pressions on the nodes. If loading of one of the expres-sions or commands fails, the whole transaction will be
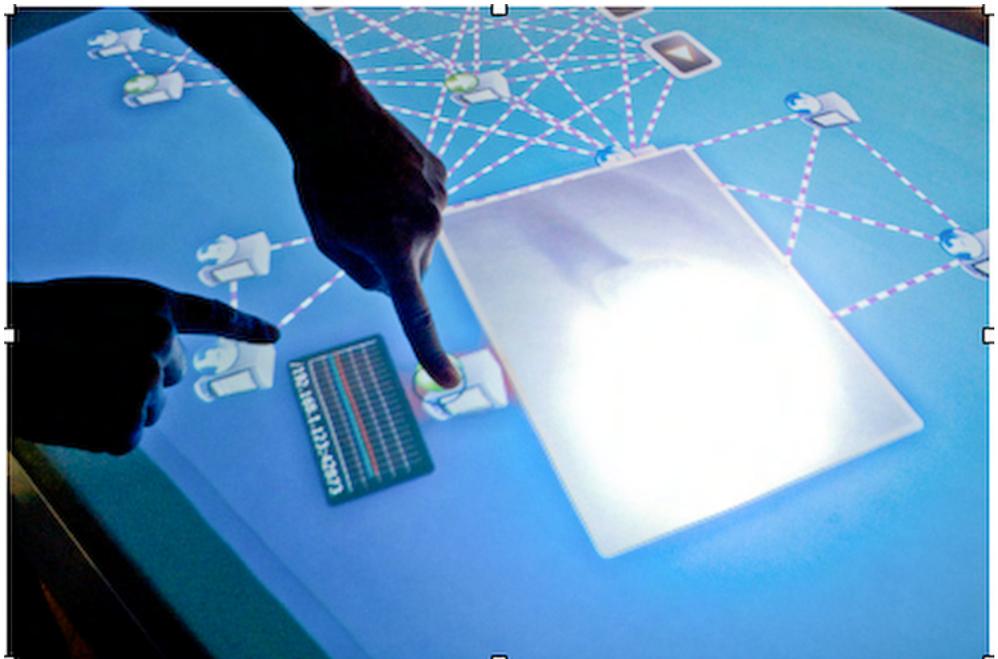
Figure 3. By touching a producer node you will see a preview of the video stream. The small graph underneath shows real-time CPU load measurements.

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

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

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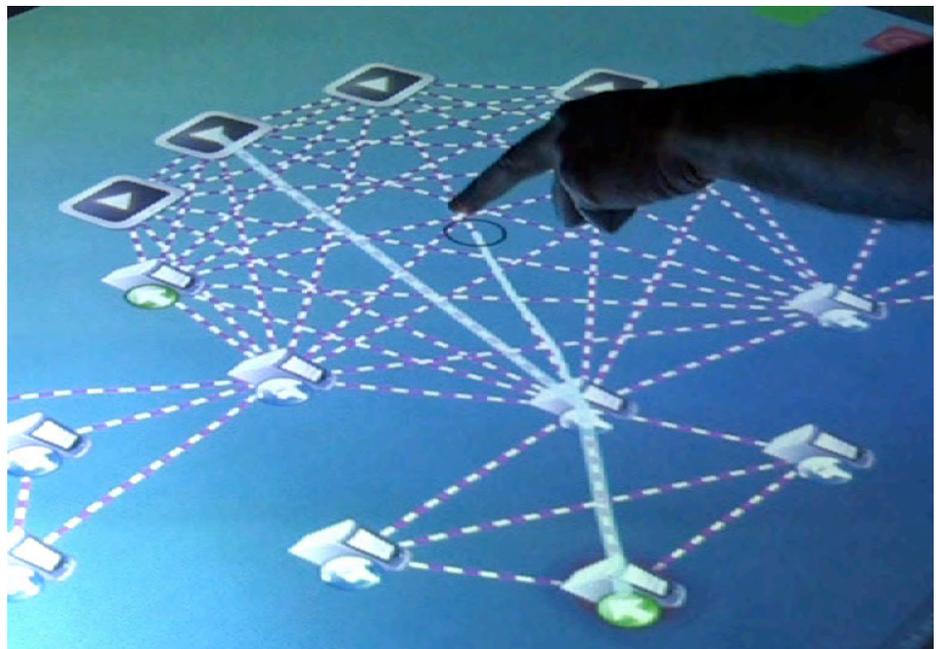
Once we have a path we can also select it and extend it to a multicast path. This can be done by dragging a new route starting from any node of the path chosen for the extension. It's as easy as that.

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?

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.

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*

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
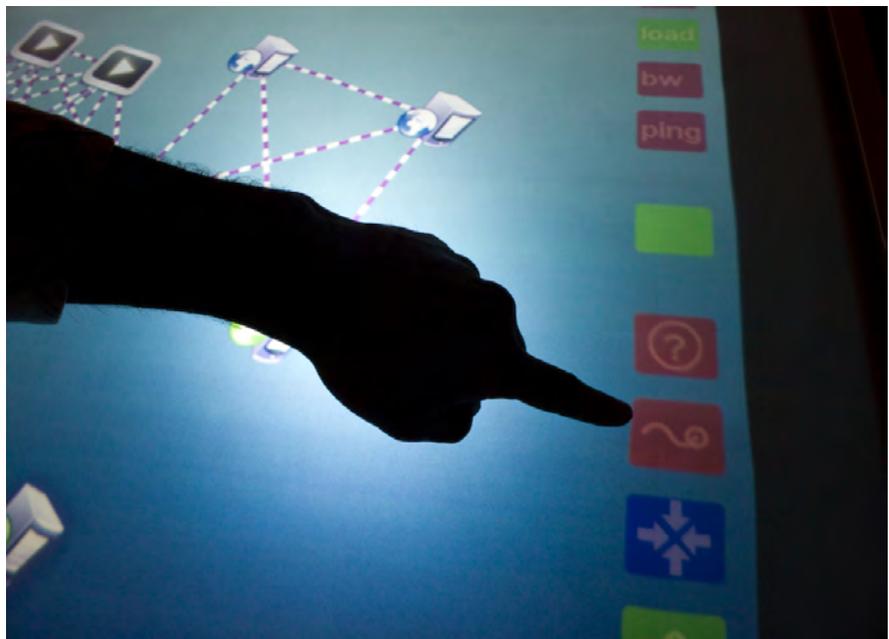
only after my pizza arrives', but also to bind network behaviour of communities or distributed applications in grid networks. Our former

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_trans-

mit 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.
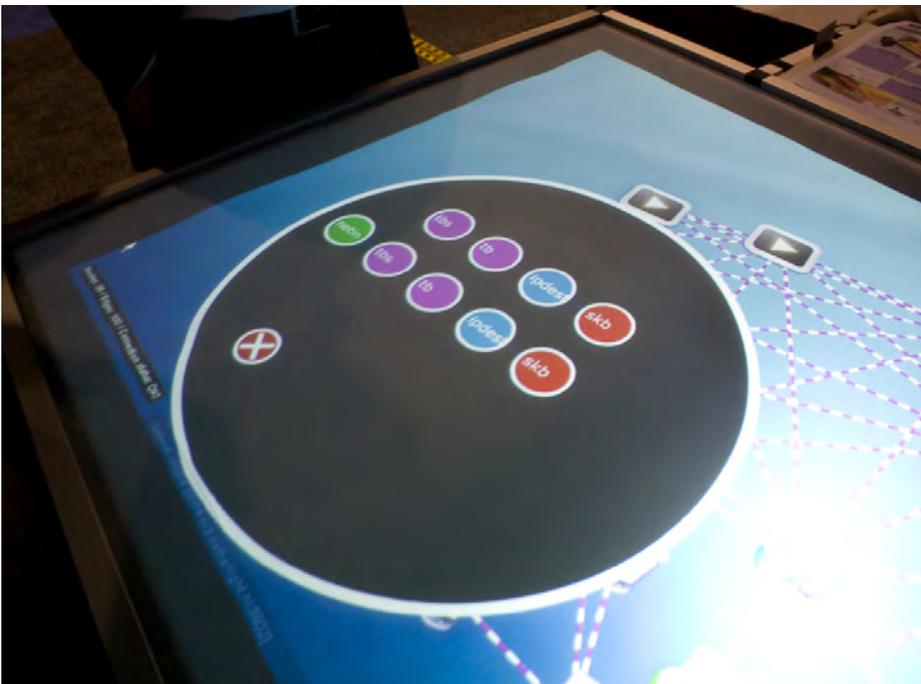
Figure 8. Streamline flow graph currently loaded in a node. Note that due to the picture contrast the connecting lines are barely visible.
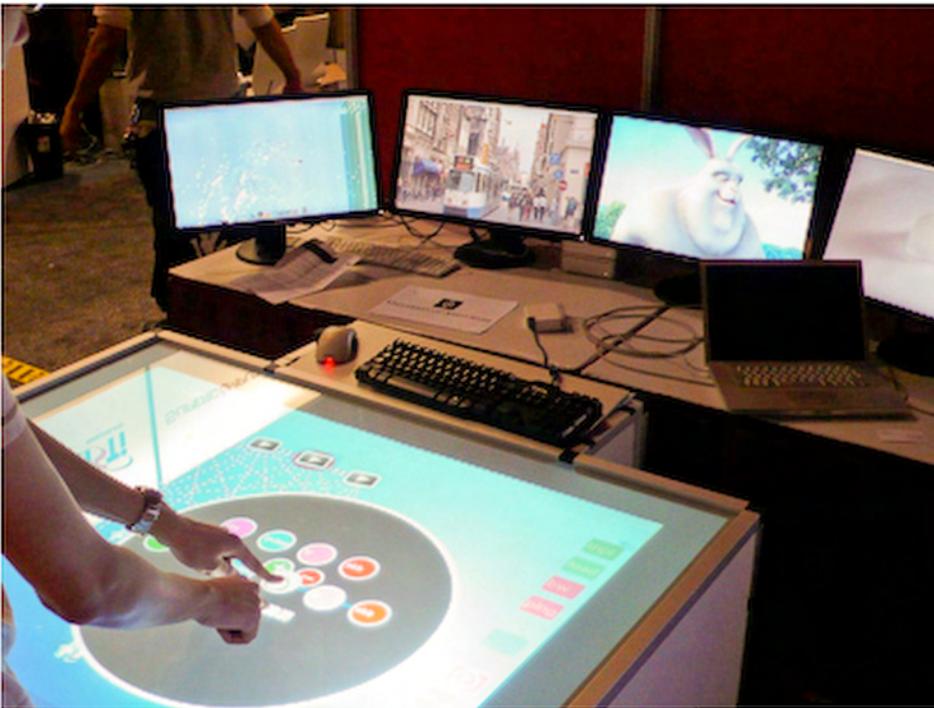
Figure 9. Modifying the sampling rate of a flow. The extreme left shows the manipulated stream, and the screen immediately to its right shows an unmodified stream that is also part of the same multicast tree.

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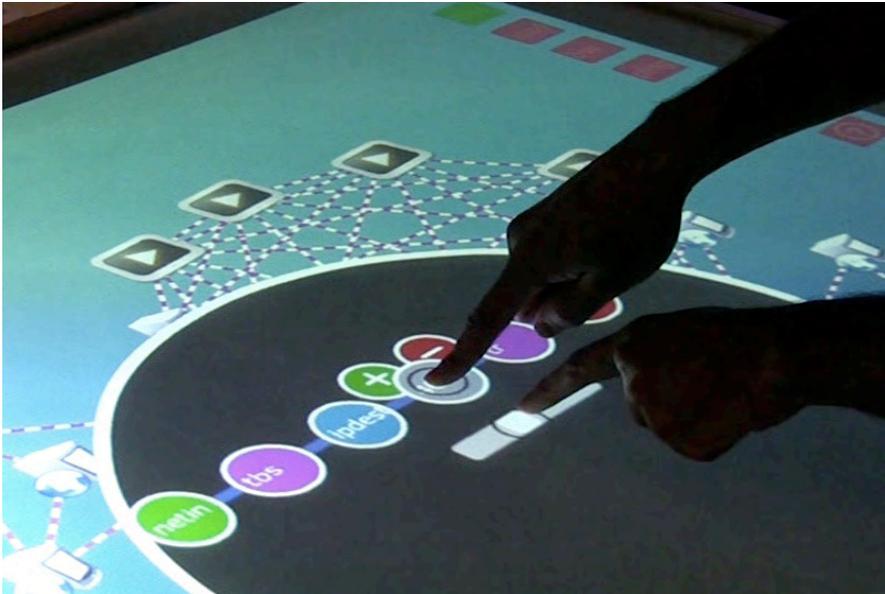
This is a powerful tool to exert very fine-grained control over traffic. For example you could add or manipulate rate limiter filters, which would allow very precise traffic shaping. The operator at the touch table could manually control the traffic shaping, but it is also possible programmatically.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

ble 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

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

measurements. And you can see what happens over there.

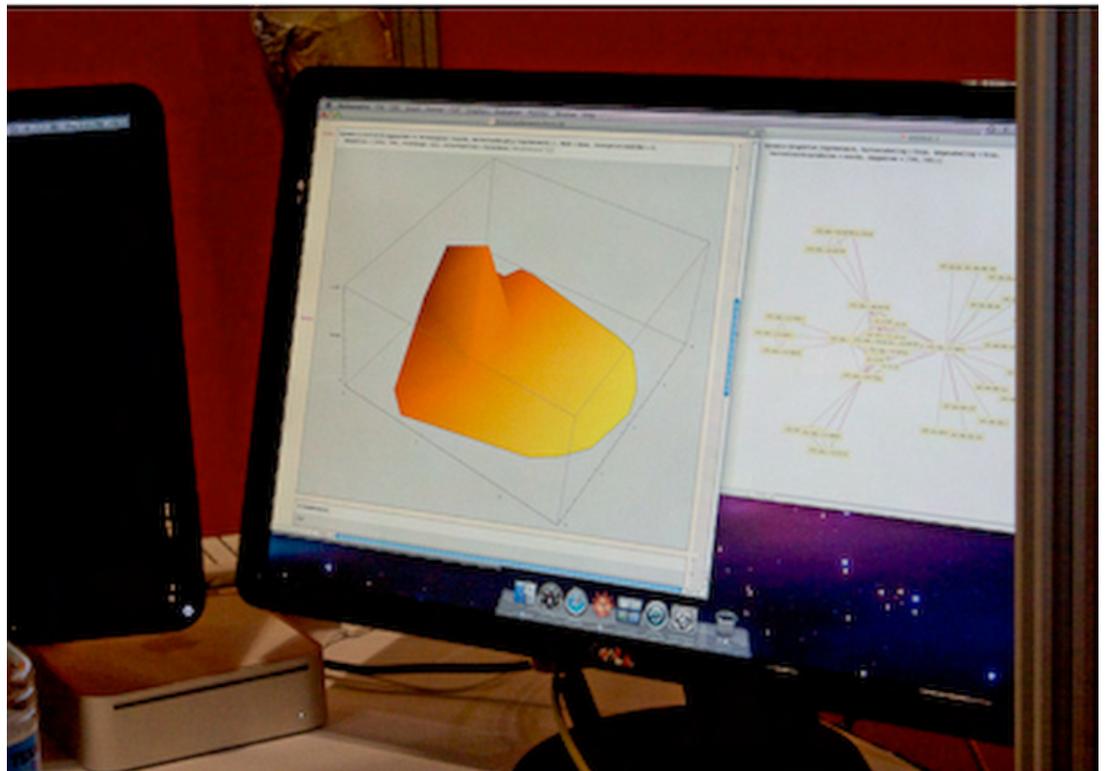**COOK Report**: You're over at Mathematica there

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 programma-

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.

and program networks inter-actively while using real-time measurements. On the touch table we do graph layouts by hand for example; Mathe-matica can calculate the lay-out 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 con-tinuous 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 net

work 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 in-formation 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 to-day illustrates a novel way to

manage programmable net-works. But, it is not limited to this case only. **Amongst other things, we would like to incorporate light-path management and in-clude 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 ap-plied in collaborative con-trol 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 net-

works 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 be-

half 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, tele-communications, 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 'life-long 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. In-

deed, 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 re-

searchers, 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 proc-

ess 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 Econo m y " 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.
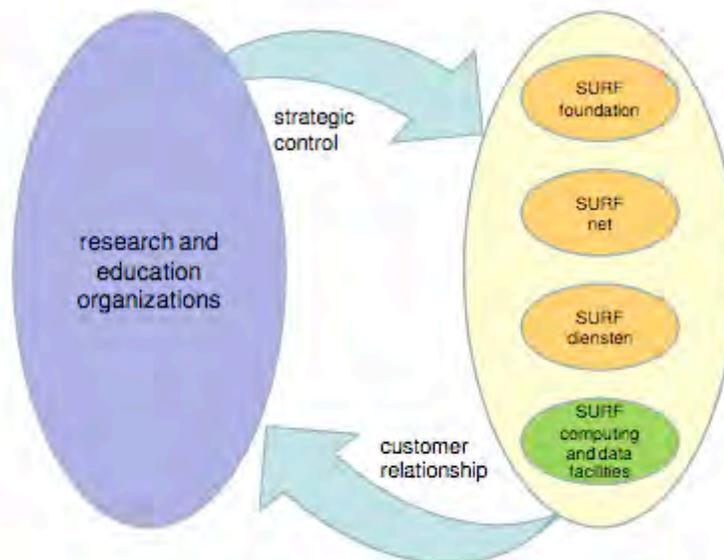

Figure 3 Proposed SURF organization

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

08 NROI 258   20/30

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 educa-

tional 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

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?

Adopting a New National Infrastructure
How Fiber Became the Basic Platform for Education, Commerce, Technology Transfer, and Research in the Netherlands

The Basis for a Future Internet?
Optical Hybrid Networks and e-Science as Platforms for Innovation and Tech Transfer

Interactive Networks

A User Interface for Application Owned Lightpath Networks

Where Is All This Headed?

Appendix: E-Science Recommendations

Executive Summary

What is the COOK Report?

The COOK Report appears 6 times a year with exclusive interviews and commentary on strategic planning for the Internet policy, economics and infrastructure. Subscribers enjoy special member privileges in the private Economics and Architecture of IP Networks mail list. A unique community of interest offering a unique online discussion where industry leaders share knowledge and develop ideas.

Membership is limited to 200 industry visionaries from commercial players such as Google, Cisco, BT, Motorola, Dell, Sprint and others. Other members represent the research and university community, supercomputing centers, network operators, industry analysts and VC communities. Also represented are community network and wireless players, both regulatory and regulatory attorneys, entrepreneurs, public sector, university researchers. 20% of member spaces are reserved for industry executives who work outside the U.S. There is a strong focus on the economic impact of open access fiber, advanced networks and other forms of rich connectivity.

Over the past 18 years special guest interviews have included Vint Cerf, JP Rangaswami, Nico Baken, Hernam Wagter, Carlotta Perez, Chris Locke, (the clue train manifesto interview) John Curran, Simon C Lin, Jerry Michalski, Harvey Newman, Bill St Arnaud and many others. *The COOK Report* has been using the mail list/forum approach to community building for 8 years. Publisher Gordon Cook was formerly served in the Congressional Office of Technology Assessment, and prior to that, the John von Neumann National Supercomputer Center at Princeton. He also has an earned Doctorate in Russian history from Duke University, 1972.

For more information visit
http://www.cookreport.com/