The COOK Report on Internet Protocol
Technology, Economics, and Policy

Building the World's Largest Optical Network to Enable Global Science Collaboration

An Examination of the Accomplishments of Harvey Newman and His Caltech Group

Editor's Introduction

Strategists should not look at the long term potential of the internet as a communications infrastructure supporting the global economy without also understanding the advances being made in optical networking by the research networks of the world. The commercial internet has been leveraged far beyond what its original designers thought possible. It is showing strains which, with the exhaustion of IPv4 address block within two years, will increase.

Two months ago we looked at Surfnet's innovative approach both to hybrid optical networks and to using such platforms to facilitate technology transfer from university research to private industry. This month and next as well we look at the innovative but, outside of academia, little known approach of Harvey Newman who has used his training in high energy physics as a platform from which to do amazing things with ultra high bandwidth optical network technology to enable extremely efficient optical VPNs as well as build a low cost telepresence system capable of global operation.

We are seeing, in these achievements, capabilities that may ameliorate some of the problems of scaling found in the first generation Internet, as well as facilitating the research needs of enterprises and enhancing the possibility of re-architecting with fiber to each desk the needs of large enterprises. We will take a detailed look at Harvey's career in this and the May issue and pick up with the June and possibly July issues on Frank Coluccio's new fiber architecture paradigm.

Volume XVIII, No. 1
April 2009
ISSN 1071 - 6327

Having met Harvey in person for the first time at Supercomputing 08 in Austin, in November, I interviewed him via Skype on December 29, 2008. By way of a general overview he sent me a slide set for a presentation he had done in Slovakia in September 2008.

With Harvey we have an interesting and rather unusual example of a scientist who understands his need for technology to enable leading edge research in high-energy

On the Inside

physics. Harvey however has not become just a user. During the past 30 years he has become a gregarious problem solver marshaling people and resources needed to achieve a global collaboration in one of the most resource intensive fields of science – high energy physics where giant particle accelerators are the tool of choice. As experiments with smaller machines filled in just a few of the blanks needed for a unified theory of physics, the exploration for new physics processes, findable only at the frontiers of ever higher energies has continued. With each new generation of experiments, and each new cycle of fresh ideas on the nature of matter and of space and time itself, came new machines of ever-increasing size. Their data output capability grew apace enabled by the progress of Moore's law, and advances in storage and communications.

With the increasing scale, complexity and cost of the experiments, global collaborations of scientists and engineers arose out of necessity, with teams based in many nations, at sites scattered across the globe. . Since the location of the machines and the people who used them were then no longer congruent with the data, networks became a necessary part of the toolset.

In this extended discussion Harvey takes us back to the beginning during the primitive pre-Web days and shows how from that time forward he marshaled cooperation to get the HEP community bandwidth and then made more effective tools for the use of the networks needed to capture and globally distribute data from for the Large Hadron Collider - the most powerful high energy accelerator and the most sophisticated set of experiments, which together represent the most complex scientific investigation ever undertaken by humankind.

Here the data needed for the science demands not only some of the world's largest optical networks working together across continental and transoceanic distances, but also tools to enhance any time, anywhere meetings among the thousands of collaborators. Especially for the past decade, after he worked out the global roadmap for delivering the LHC data to its researchers, Harvey has been roaming the world conducting the problem solving approaches needed, and building tool sets which are potentially usable by many other fields of "data intensive science", including fields such as climatology, fusion energy and bioinformatics whose needs are also about to "take off".

He has the vision to see and express possibilities on a grand scale. This piece came about by chance when I sat next to Rogerio Iope from Sao Paulo Brazil on the plane flight from Newark to Austin Texas. Rogerio [shown below] is a doctoral candidate who was going to the meeting on a budget, and was brimming over with enthusiasm that Harvey had invited him to work with a group of about 12 others at the Caltech booth on the show floor where they were preparing for and executing the 2008 Bandwidth Challenge. Although I cannot claim to know Harvey well yet, Rogerio's spontaneously expressed enthusiasm for working with this man was infectious. I shared my room with Rogerio and consequently got daily reports of what was going on at the Caltech booth and on the next to

the last night, after the show floor closed I dropped by the Caltech booth where Harvey was coordinating a team of about a dozen people who were working at a long table along with ten others who were coming in by video links from around the world. Harvey spent the next two hours explaining what was happening. I was fascinated. My initial interest that was later cemented by Harvey's previously published late November comment to the arch-econ list:

"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."

This interview is the result of those experiences. I remain convinced that together with the Netherlands' work on hybrid optical networks, there are public policy lessons here focused on the need to invest in national optical infrastructure in ways that will advance science education and help the United States remain economically viable in a global society.

Consequently this and the following issue present an extensive introduction to the high energy physics community's approach to global networking and science education.

The Servers at CalTech's Booth at Supercomputing 08 on November 19, Austin Texas

Why Networks?

COOK Report: Would you give me a quick overview and summary of how you began to get involved in the use of networks for your research in high-energy physics? I understand that by 1982 you were a Professor of Physics at Caltech involved in research in high-energy physics with particle accelerators and becoming a part of a global community of high-energy physicists with a need to find the best means of sharing the and understanding the output that these scarce and rather expensive machines produce.

Newman: In the beginning our getting involved in networking was motivated by the challenges we faced when for the first time a relatively large group of physicists based at US universities was going to work overseas. We were going to DESY (the Deutsches Elektronen Synchrotron) in Hamburg which had the highest available energies for the kinds of investigations we were doing. And the question in my mind was how could one be working over there but based in the US? How can you do this and be effective? And how do you make the tie between the lab and all the young students and postdocs who cannot always be at the laboratory?

By 1981 I had already been involved in these experiments for some time. Actually the first time I worked overseas was in 1974 when I started to work at CERN. However at that time, through the late 1970's, all the activity was focused locally. In 1981 the key event for me was that after six or seven years in Europe my children were about to start school and for that reason I had to think about going home and being based in the US. Obviously I then began to wonder how you could work with this center that was based in Hamburg? Despite the fact that I traveled frequently and spent summers there, this was clearly not enough. Consequently I began to investigate networking very seriously.

COOK Report: What was the cause of a large group of American physicists going to work at DESY in the early 1980s?

Newman: In the late 1970s and early 1980s we had already seen the situation change to the point where a lot of the leading edge research was now being done in Europe. For my thesis experiment, I had worked at the Cambridge Electron Accelerator at Harvard, where one of the first "colliders" had been built and which produced head-on collisions between electrons and positrons. After getting my Ph. D in 1974 I spent a few years at CERN in Geneva where I worked at the so-called Intersecting Storage Rings on experiments involving proton-proton collisions.

By the late 1970's it was clear that the most exciting experimental programs was going to be at DESY in Hamburg, where they had built the highest energy electron-positron collider, conceptually similar to the one at Harvard but with an order of magnitude greater energy. So I moved to DESY in 1978, employed as a staff member there.

In 1981 when I first took a job on the faculty at Caltech while continuing my work at DESY and beginning to plan a future large experiment at CERN, I began to wonder how one could possibly run a research team and supervise students based in the US, while being deeply involved in laboratories overseas. This motivated me to begin my work on networks, and indeed to found the whole area of international networking for high energy physics. Now 26 years later I continue to work and make progress in networking for science, in addition to my primary work on fundamental physics.

COOK Report: Were the Europeans at this time beginning to invest more in this kind of research than the Americans?

Newman: It was the beginning of a long-term trend. Remember that in the 80s you still had Fermilab. They had built a machine that would collide a high energy intense beam of protons with a target in a way in which you could get a high number of interactions but with a lower center of mass energy – the center of mass energy being the framework within which you can create new things. But Fermilab would only produce a proton-proton collider in the mid-to-late 1980's.

When I moved back to the US in 1981 my interests – then centered at DESY -- were focused on the electron positron collider area, and we had just designed a new much larger experiment that was planned to run at CERN, where they were building a new electron-positron collider called LEP that would reach an energy more than four times higher than the one at DESY. So I decided to continue my work at DESY through the mid-1980's, while also traveling more and more frequently to CERN where our new experiment called L3, was being built.

Evolution of Network Infrastructure in the 80s

COOK Report: Would you take me then through the evolution of network infrastructure from the early 80s onward into the 90s?

Newman: In 1981 the first international networks we used were based on the OSI protocol X.25. There were some X.25 services provided by companies in the US like GTE/Telenet to which we could subscribe. There was also a service provided by Deutsche telecom called Datex-P (P for packet).

COOK Report: The cost of moving significant amounts of data by that protocol would have been prohibitive would it not ?

Newman: It was volume sensitive but as long as you didn't do too much it was not too steep. I could send you a bandwidth roadmap of mine that is not in the deck you're looking at that starts at 9.6 kbits per second (kbps). But that roadmap started in 1985 with our first transatlantic "leased line"; in 1982 we were working at 2.4 kbps, less than one-ten millionth the throughput we deal with today. For what it was, it was indeed expensive.

When I was at DESY, and began to do the first trials with networks to the US, we always had to rely on those first "packet" commercial services. They were all there was. But even then we had to do development to make them work. We found there were some parameters in the set-up packets (X.25 is connection oriented and hence it was efficient for very low bandwidth links with volume sensitive charging) that Telenet had screwed up; it was sending US standard parameter values to the international service, which had different parameter positions in the packet. So it was clear that the service could not work internationally as it was doing within the US.

But we then used some of the X.25 debugging software utilities on our "big" Digital Equipment Corporation VAX computer (which had all of 1 MIP of computing power, but consumed tens of kilowatts) and using these to examine the packets being sent, we discovered the mismatch and told Telenet how to fix the problem.

In the early 80s we put together a small set of X.25 packet switches, of the same type then used by the British Post Office. These were prepared for use, following a configuration I had worked out, with a few of my fellow physicists who were working

on the experiment that was then being prepared for CERN, from Michigan, Princeton and DESY. They started at 19.2 kbps and, after a couple of years, went to digital connections at the grand speed of 64 kbps per second.

By 1984 we realized that - since these were volume sensitive charging services - some people looking at trend lines would become very concerned about what would happen. And once or twice we got some pretty big bills at the end of the month. As a result we began looking into leased lines in 1984. Our first leased line was between Geneva and Caltech. It including microwave links between Geneva and Vienna, and was provided by "Radio Suisse" at a cost of more than 100,000 Swiss francs for a 9.6 kbps connection. But at least it was not volume sensitive.

In the mid 1980s I was involved with the technical advisory group for the emerging NSFnet. During this time we met and began to work with people who had spent their lives on TCP/IP. Consequently, given our X.25 experience, we ran a multi protocol network for a number of years. Our 9.6 kbps link became 19.2 kbps as we got better, more capable (and expensive) modems and eventually moved to 64 kbps. We were running DECnet, TCP/IP and X.25. The struggle was one of what would run over what? Would it be Decnet over IP? Would it be IP over DECnet? Or DECent over X.25? I was probably the only one in that group who really understood what happened to the IP packets on such low bandwidth networks, versus what happened when you were working with other protocols like DECnet/X.25, or IP/X.25 for that matter.

The TCP headers were large for the networks of those days, and so as people would use character transmission on terminals, TCP would often send a packet of 56 bytes for a character. So the results were not necessarily a given. Now I should also mention that we had BitNet store-and-forward transmission of email.

Under these conditions the NSFnet would have been transformative except for one thing. It was a 56 kbps network and by 1988 or 89 it was dead; it was simply overwhelmed by the demand. When with Hans Werner Braun they began the transition to T-1 in the US, it didn't help me too much because my needs were international.

But by the end of the 1980s we did get speeds of up to 256 kb per second running on our multi protocol network. It had had a larger and larger proportion of TCP/IP as its traffic at this point. Around 1988 - 1989 the American Physical Society had a magazine called *Computers in Physics*. I was at an editorial board meeting in New York City and I remember that people were groaning about how the NSFNET network was now completely useless. At this point the transformative thing for us was a relationship with IBM.

Ushered into the 1990s by Al Weis and IBM

By 1989 IBM was getting involved in international networking for research, and Al Weis, one their VPs in this area was thinking of going to DESY. I knew Al through our contacts on a National Academy panel on "Information Technology and the Conduct of Research" in 1986-7, and through other contacts in IBM, and so I said to him: No. You should go to CERN where we were about to start our large scale electron positron experiments.

We had a rather substantial history by then in networking, were interested in greater use of TCP/IP, and we had some lively conversations which led to his visiting CERN. After a single day of the visit in 1989, which I hosted together with the late David Williams, who was then

the head of the Computing and Networks division of CERN, Weis changed his mind about DESY and decided to fund a T-1 link from New York to CERN -- a link costing IBM $3 million a year. He funded that from 1990 to 1993. I think that during that brief period he had become convinced that CERN was the best place for him to support his interest in the use of networks for research in science and academia.

What I think we did do was wean him away from his feelings that the only place where people did serious networking were supercomputer centers, and perhaps helped him to see that the potential for networking in science and academia lay elsewhere.

COOK Report: The timing involved in your showing him what TCP/IP can do is interesting. I arrived at the Office of Technology Assessment in Washington the week after Labor Day of 1990, and within a few days of that, the NSF and Weis announced an amendment to the Cooperative Agreement that formed ANS (Advanced Network and Services). IBM committed to work with MCI to increase the the T1 backbone to T3 – all admirable goals -- the result of the change was indeed a better network for NSF's constituency but it was also a test bed for IBM's development of a TCP/IP router that could enable its mainframes to connect to networks using protocols other than IBM's proprietary SNA and VMS. Previous to this they required extremely expensive front end processors to interconnect their mainframes.

Just an interesting note of policy history and, although I did not like it at the time, going on 20 years later I have to say that I think it was likely not a bad deal for all concerned

Newman: It is true that networking on IBM mainframes throughout the 1980s had not really been focused on TCP/IP. They were indeed focused on SNA which had roots in the OSI model.

But going back to the timeline I was telling you about: we ran our own multiprotocol network through 1994. Now I had written to DOE as early as 1985 asking them to support some of this. And then in 1995 I established a consortium with CERN into which DOE supplied some funding. Caltech representing DOE and the US high energy physics community, and CERN and the French computer center in Lyon representing the French academic network and WHO at the UN in Geneva, joined together in a consortium which bought a 2 Mbps link -- known as an E1 -- the European equivalent of a T-1. Over a two-year period (I996-1998) we were able to go from 2 Mb to 4 Mb per second during a period of time in which there was still regulation for carriers in Europe, and during which the only links available were E1's.

Deregulation Opened the Floodgates

Starting in 1998 we were greatly assisted by deregulation. This meant that instead of having Swisscom and maybe one other vendor bid, within a year after deregulation we ended up with about 12 different offers for our data network from various carriers. The second important event of 1998 was that vendors were allowed to sell other products than 2 Mbps E1 links. As a result of this we bought an ATM link. It started at 12 Mb per second in 1999 and went to 20 Mb per second in 2000. Our first optical link was an STM-1 (155 Mbits per second) which we acquired in 2001.

We then started to go up the optical steps moving to 622 Mbps in 2002, 2.5 gigabits/sec (Gbps) in 2003 and in 2005 10 Gbps. We then went from two to four 10 Gb links that we have now and we have a roadmap to go to eight 10 Gb links by 2010 followed by an aggressive transition to a combination of

40 Gb and 100 Gb. In other words we intend to start at 80 Gbps across the Atlantic in 2010 and increase that total by 1.5 every year until we reach 400 Gb per second in 2014.

COOK Report: How do you envision that transition process occurring?

Newman: Before the transition we will be starting with multiple links following multiple paths (to ensure non-stop operation of the network even in the presence of one or two cable cuts) and eight 10 Gbps waves. We will then add a 40 Gbps wave giving a total of 120 Gbps. In the next step we will replace the eight 10 Gb links with a 40 Gbps per link and our first 100 Gbps link, and then subsequently move to four 100 Gbps links across the Atlantic by around 2014.

When you look at this roadmap and stare at it for a while you realize several things. Back in the 1980s and the 90s we used to be going through technology generations of the physical means of data transmission. Now we're not doing that anymore. This means a slower evolution. And therefore my plan is to only increase by a factor of 10 over a period of five years between now and 2014. This rate of evolution, which is equivalent to a factor of 100 in a decade, is much slower than in the past. If you look at the roadmap between 1995 and 2005 you see a rate increase of 5000 times in a decade, and a factor of 1 million increase in bandwidth over the 20 year period 1985 - 2005.

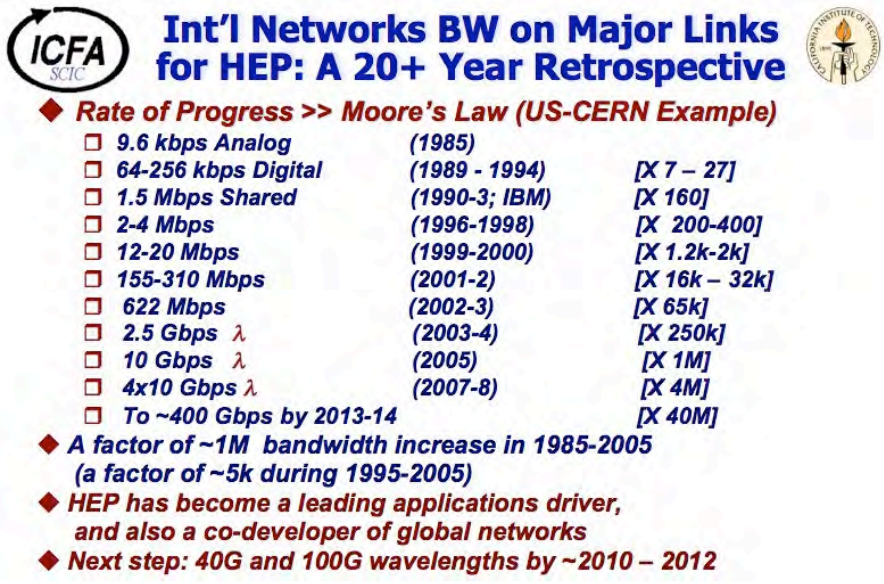

The increase has thus been about 1000 times per decade or about doubling every year on the average.

COOK Report: When you say we are no longer changing technologies to what precisely are you referring?

Newman: Up to the mid-80s we were dealing in analog links from the mid-80s to the mid-90s in electronic digital links - 64 kb per second and in multiples thereof – then T1 to T3 were still electronic links, although of the next generation. We made our first use of optical links in 2001 with STM 1 and the OC 3s and then changed to 10 Gbps by 2005. We gained these benefits by changing not only the transmission protocol but also the network interfaces on the switches in the network.

What has made the wide distribution of data to and from many switches and also many servers at one and now 10 Gbps feasible was the abandonment of the old "carrier class routers" and the adoption of 10 gigabit Ethernet (10GE), in full-scale switches where a 10GE port costs on the order $ 10k. We had a Juniper router in service for a couple of years and it was a real stretch for us to finance a configuration with just one wide area (OC-192) port, one

10GE port, plus a few one GE ports.

Using Force10 switches supporting 10GE WAN-PHY, Cisco switches with many additional 10GE ports, and now CIENA optical multiplexers where the cost per OC-192 port is similarly low, we are able to continue our upward bandwidth migration. We now have dozens of 10 GE ports in midrange and low-end switch-routers. **But no carrier-class routers.**

We are going now in directions where (with the use of large optical displays in e-science applications and in basic fiber to the home and in the increasingly commonplace appearance of gigabit ports in consumer machines) demand for 100 gig Ethernet (100GE) is increasingly strong. Because of rising demand, typically 60% Compound Annual Growth Rate (CAGR) in the developed world and 80-100+% CAGR in the developing world, vendors are not only hungry for implementation of 100 Gbps transmission standards but even for pre-standard implementation.

The good news is that for the first time, as I heard in last fall's Internet2 meeting, that the 'telecommunications people' and the 'data communications people' are actually talking to each other. Consequently you have a consistent set of standards beginning to emerge for the OTU hierarchy that accommodates both SONET and Ethernet.

COOK Report: SONET is a telco standard. Is it going away or is the installed base so large that we are stuck with it?

Newman: Commercial services are OC-768 which is 40 gigabit SONET. But then there is this OTU hierarchy. I'm not sure we are talking about SONET any more. The same sort of hierarchy however will fit into the new OTU standard so whether you have SONET or Ethernet they are simply accommodated into the larger framework. When we did our press release in November with Ciena we mentioned that we had a real OTU4 single-wave 112 Gbps link with a 100 Gbps payload. The OTU standards will accommodate 100 gigabit frames no matter whether they are in the SONET hierarchy or Ethernet.

Let's talk about what is possible. The world now is focused on this transition to 100 Gbps and for this reason they are not thinking much about what lies beyond. Nevertheless I think there's good reason to believe that in the second half of the next decade we will have terabit per second links.

Another thing that can happen is the commoditization of wavelengths, just as now gigabit per second Ethernet interfaces cost essentially nothing, and also in servers even 10 Gb per second Ethernet chips on the motherboards cost very little. Within such a context, one can imagine future systems with many wavelengths transporting many Terabits per second.

COOK Report: But getting these many many wavelengths across long distances by carriers is no easy proposition.

No Progress Until the Telco Model Becomes Technology Driven

Newman: True. The world is politically and policy driven not technology driven. Somebody mentioned today on your list that the carriers seem to be waiting to see public demand. However this is not true for storage and not true for processors in that the makers of these devices just keep increasing their capacity. **Unfortunately the deployment of bandwidth is not technology driven. The vendors are looking at markets.**

We seem to have a kind of conundrum with bandwidth: you do not seem to have an elastic demand.

If the producers put out 10 times the capacity, the public seems to be unaware that it's there and unsure of exactly what they should do with it. With bandwidth neither awareness nor a real cost model is driving an elastic demand.

With Intel's new Core i7 chip the most recent model is actually a bit faster and costs the same price as the previous low-end model. It is therefore enticing to transition to. People don't have to pay any more money and they are curious to see what it would be like, so they buy it.

To achieve this with wide are data transport, the vendor's need is not only to build a model of an elastic demand but also one with the commoditization of switches and interfaces as well as the commoditization of 10 Gbps links. This is coming but only in servers not yet in PCs on the mass market. So some parts of the information technology world are technology driven and other sectors including the deployment of bandwidth are demand and market driven.

COOK Report: And when you get into that universe you are stuck with the chicken and the egg problem.

Newman: As well as in the deployment of new technologies where there is no concept of Enablement to explain what we can do with new resources.

COOK Report: Do you have any ideas on how to overcome these problems?

Newman: One idea would be to use it for education and research. You know just give it away because then the young would grow up with a different point of view and create then an elastic demand model, because they would be used to using the latest technologies to innovate and create.

COOK Report: The idea in my mind after the supercomputing show in Austin is an opti-portal in every school.

Newman: Sure why not? But **the problem in the United States is that there are no exceptions to the usual business models for research and education. The only way to give benefits to that community is through private nonprofit corporations (such as Internet2) being formed with the mission to help the community. The problem is that there is no place that they can get beneficial pricing. They may get lower pricing but it doesn't mean they can deploy a factor of 10 or 100 over a period of several years and get it all out there because there is no technology driven model for them either.** The market does not give them that benefit.

COOK Report: But these seem to be choices that we are have to make at the national strategic and political planning level. They are not the kind of choices being made in countries like the Netherlands. We seriously need to change our strategy direction.

Newman: Right. It's a different strategy. **And that's what we discussed in the "highway model" based on dark fiber infrastructures which have been deployed by national research and education networks (NRENs) in a growing number of countries outside the US, including economies like those in central Europe (Czech Republic, Slovakia and Romania for example) that are orders of magnitude smaller than the US. Just put things out there. And find out what people want to use. Create conditions where the costs are low so that mass-produced technologies to use the infrastructure fully, along with the methods to use it effectively on a large scale, are created rapidly, and**

evolve as the infrastructure evolves.

COOK Report: I think we need a chain of reasoning or a roadmap, were it possible to have such in the United States, that could be used for discussion and potentially a change in our direction. Why can't we change the agenda in education from what you said of being one where only private entities can do things to the point of view, of one where for the sake of American competitiveness, we change how we allocate resources?

Newman: A few years ago I was discussing strategies with a person prominent in this field and in a position to make an impact in this area. **I said that we really needed to discuss strategies for getting new technology into the hands of the R&E community. What he wound up having to explain to me was that you were not allowed to give away what a company can sell you.** Apparently this is written into law. You can give something you already own or possess to your own organization, but if you are the government, you cannot give something away in the private sector that someone else could sell.

COOK Report: Somehow that has to be changed.

Newman: Well I will give you another perspective. When we were building up this road map there was a time when I thought having a close relationship with an (unnamed) incumbent would help us – and it does a little bit. I remember being in conversations on this with people at DOE and the lady from the incumbent was in on the call and she said why don't you let industry do this [referring to our developments and roadmap building and the long range vision that goes with it] ? So I replied "Okay … are you ready to step up to the plate ? Will *you* do it?" Not surprisingly there was silence.

The incumbent telcos don't develop uses. And there is great benefit in giving things like this to us so that we can show what can be done with it and give other people ideas about what they can do. We can create new business models for them. But they simply don't think this way. They simply do not think in terms of exponential growth in capacity that generates more demand for still more capacity and yields more revenue.

COOK Report: I think for some time now that this is a unique American problem; that it does not extend so much into other countries. Where the solution will come from remains to be seen, but in the meantime could we use it as a transition to the events that helped you become involved in bringing the network for the LHC into other countries?

The New LHC Network

Newman: We went to DOE for the first time in 1996 when there were all these plans going on for US involvement in the LHC project. Plans for the Texas-based superconducting super collider ended in 1993 and by 94 we had joined the LHC program. There was a lot of planning in 94 and 95 for building the LHC apparatus but it was only in 1996 that the question of computing came up – since we raised it.

I made a big drawing of a large computing Center at CERN and another in Fermilab in the US and I began to talk about it to the NSF and to the DOE. One of DOE's first questions was "… you're not asking us to put a lot of money into a big computer at CERN are you? I looked at them and said "No, I would never be so foolish as to ask you that."

I got the message, and the message was really universal. It was that if you want to gather resources for all these groups working together, then some of the resources you gather for them needed

to be local to them – within their country, within their region, and even better at their university or lab. You could not take the resources and send them all overseas and have that draw the focus of attention entirely away from the universities from which those resources came. You needed to bring the vitality of scientific discovery home to your colleagues and students; not just drain it away to some far flung place on another continent.

So in 1996-8 we developed the concept of a worldwide distributed system for LHC computing. We soon found out that it shared many of the concepts of "The Grid" which had just been developed by Ian Foster and other colleagues in computer science, except the high energy physics case would have to be much more "data intensive" given the massive data volumes the LHC experiments would have to process, distribute and analyze. The high energy case would have to be much more "stateful" since computation and data would have to be brought into coincidence. It would have to be more dependent on high performance networks since large datasets would have to be moved quickly from place to place to meet the needs of the scientists to analyze the data (using distributed computing resources) once it had been processed at one or more of the larger centers.

The ability to satisfy local and regional and national resources drove this. **Our first proposals to NSF emphasized bringing the intellectual strength to the universities and not draining it away. This got very strong support in and built upon the idea on which we had been working for many years – namely that networking was central to success in carrying on an ongoing global collaboration.**

In 1999 I was called by NSF for a key meeting, where they wanted to know how they (NSF) could get involved in, and establish a special role in computing for the big projects they were funding, including the LHC and also LIGO - the Laser Interferometer Gravitational Wave Observatory. In order to indeed bring home some of the resources and intellectual focus of the distributed computing model to the universities, I then introduced a hierachical picture of facilities of varying sizes, from the "Tier0" central facility at

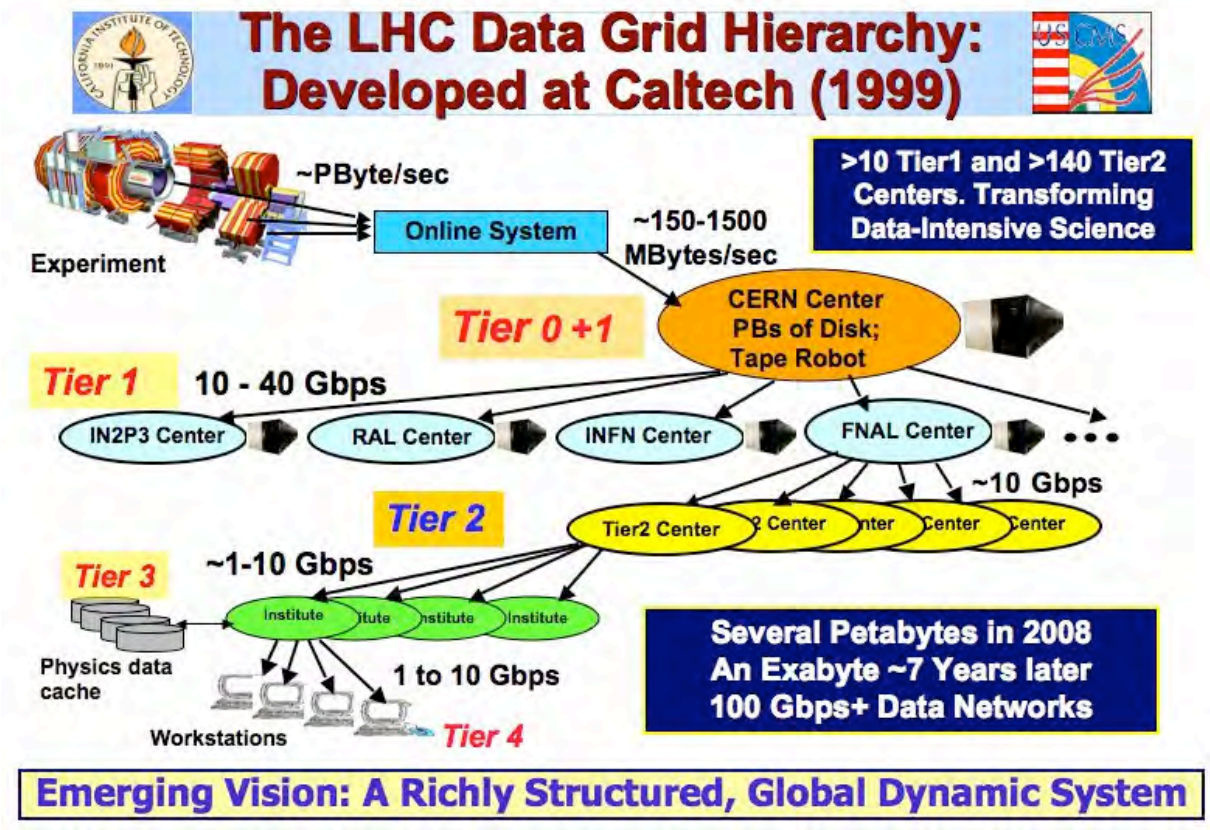

CERN, to the large national "Tier1" computing and storage centers, to "Tier2" centers of a few racks of servers and storage on a scale that could be built and managed at a mid-sized or larger university together with adequate networking, to "Tier3" clusters serving each individual physics team.

It was the Tier2 center (the first of which I designed and built at Caltech in 2000) that turned out to be the key, and which made the whole concept compelling. If you have a several rack facility of computing, storage and networking, then you can draw some attention to your university locally. Indeed some of the less-rich countries also became "Tier2", from Brazil to India to Estonia -- and now 140 others. They understood that they could draw attention of local groups to their facilities because they were **locally a player but also, because of the network, one on the world stage**. This idea had tremendous traction.

The original hierarchical drawing of what the worldwide computing model should look like brought the whole concept into focus in everyone's mind. It became iconic. The hierarchy has since stayed the same, but as bandwidth capabilities have become greater and greater I have changed the capability of each level in the hierarchy to match the present reality over time. And while we know that the actual data flows are more complex that the simple drawing shows (for example a Tier2 may have to draw data from many Tier1s) "the hierarchy" concept is still the basis, and the university-based Tier2 is still a vital part.

Drinking from the Collider Firehose

[**Note to readers**: the remainder of this installment describes the collider and its research goals - fascinating material but if one is reading ONLY for networking information please turn to page 30.]

COOK Report: Is it correct that the output of the collider is as much as one petabyte per second?

Newman: That's just what comes out of the front-end readout systems on the detectors. The detectors and real-time systems put out somewhere on the order of between a terabyte and a petabyte of data per second. But then we do a selection process on the data which results in storing only a few hundred megabytes of data per second. This is done through a multi-step multi-algorithm filtering system where you have different selection criteria for each of the large number of physics reactions. You must make sure that between all the physics goals of the experiment you cover all the bases with your trigger algorithms. So even after all the online filtering, you end up storing a very large number of events, and each event, representing one head-on crossing of a pair of proton bunches in the machine, is very complex.

To tell you some of the numbers - at the design luminosity, that will only be reached after about a couple of years of operation of the machine, you would have on the order of 20 interactions for every crossing of the bunches of protons, and since there are 40 million bunch crossings per second, you end up with close to 1 billion interactions per second. Of those billion interactions you keep only a few hundred.

COOK Report: Then a huge amount of filtering is done immediately at CERN?

Newman: Yes. There is an online filtering process where the lowest level is done in dedicated hardware and then the higher levels of the filtering process are also done using thousands of processors. And after that only one part in 10 to the sixth or 10 to the seventh of data generated is actually written to disk for later analysis.

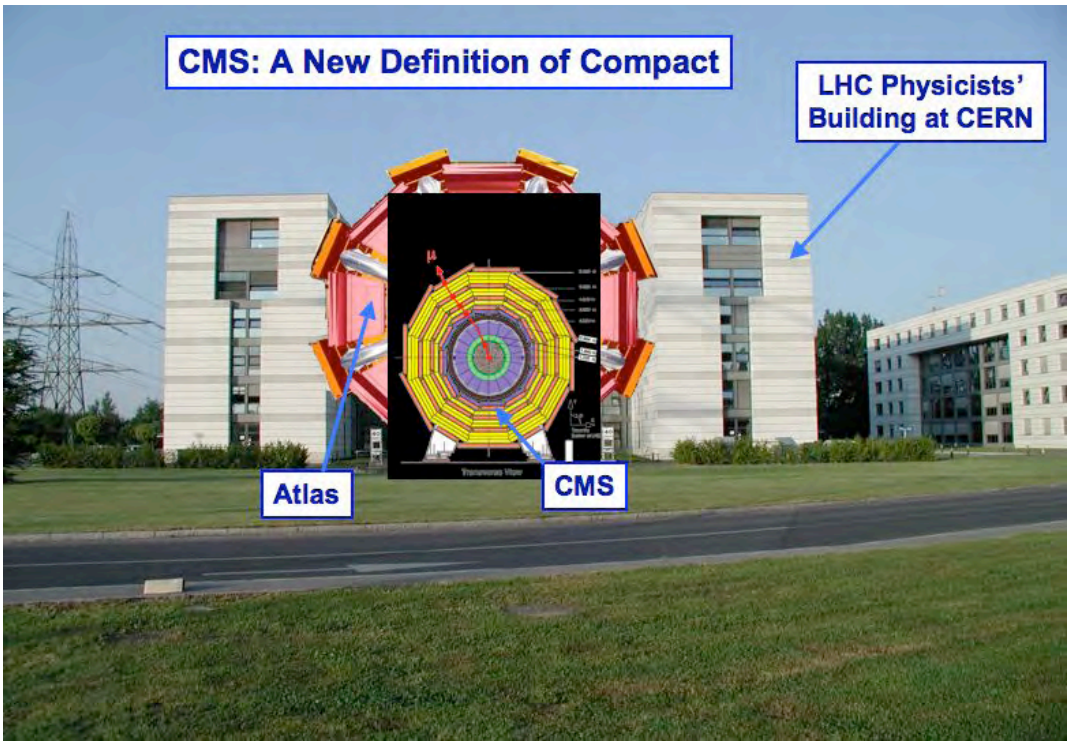

You have colliding protons. Now a proton is a complex system with three quarks bound by a field that also appears in the form of other particles called gluons. The interactions can be interactions of any energy. We actually call it momentum transfer, where momentum is transferred from one constituent in one proton (a quark or gluon) to another constituent in the other proton in the collision. We are only really interested in the very high momentum transfers or more crudely speaking in the very high energy collisions between the constituent parts of the proton, because we already understand from our standard model of particle physics how to describe the lower energy interactions. Therefore in order to search for new physics processes, we select only the high energy interactions.

COOK Report: Would you take me through the background setting of those first slides.

What We Are Doing with the Collider

Newman: The slide immediately to the left above shows the collider tunnel and an aerial view of its setting on the bor-

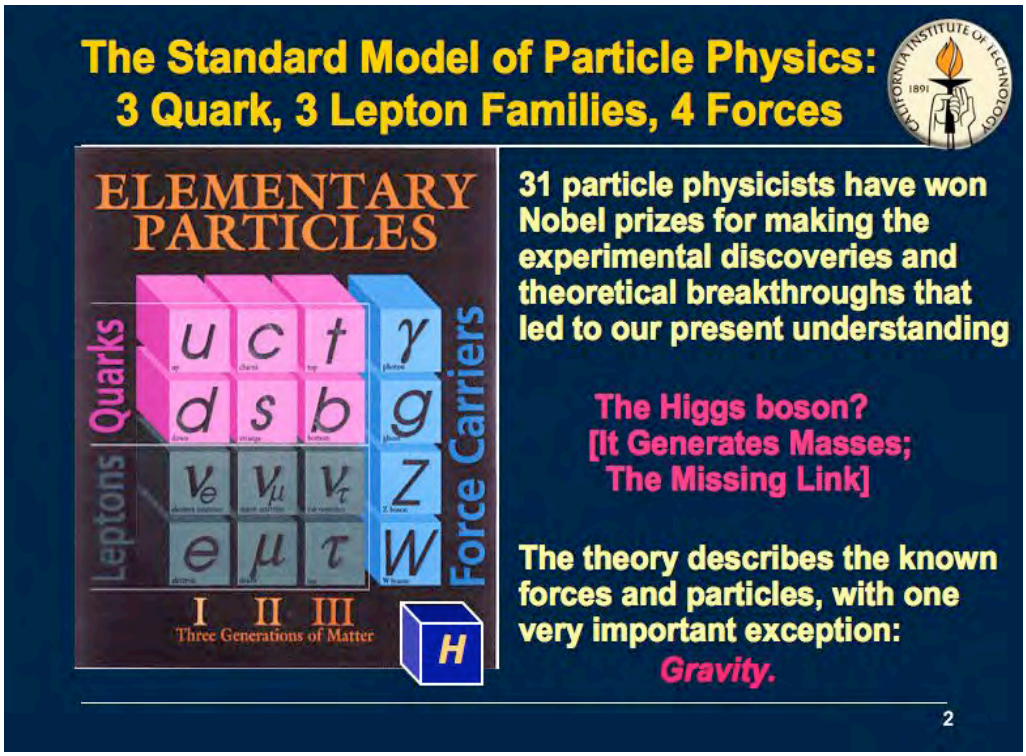

der of Switzerland and France, not far from the shores of Lake Geneva.

The second slide on page 14 above shows the Atlas detector superimposed on the physicist's office buildings at CERN. It is taller than the six story building. The ATLAS detector contains a series of ever-larger concentric cylinders around the central interaction point where the LHC's proton beams collide.

And the the "Standard Model" slide shown above is a beautifully simple picture of the fundamental constituents of matter and their interactions.

There are three families of quarks, three of leptons and four forces. Normal matter is made up of the first family. Consequently Protons and neutrons are made up of up and down quarks. The electron together with nuclei is what makes up atoms.

COOK Report: And the force carriers?

Newman: The strong interaction gluon, the "g" there, binds nuclei together. It is a strong short-range force. That's what it does. The photon is the particle that carries electromagnetism. The "Z" and "W' particles carry the weak interaction which is responsible for radioactivity. Another part of the standard model is unification so the weak interactions and the electromagnetic interactions are shown to be two aspects of one unified interaction called the electro-weak interaction.

This is a tremendously simple picture for describing all that it does. But it has a few problems, one of which is that it does not describe gravity. The way these theories are expressed is quite analogous between electromagnetism, the "strong" interaction (which at high energies actually becomes weak-looking through a process known as asymptotic freedom, which allows us to describe it theoretically in a way quite analogous to the electromagnetic interaction) and the weak interaction.

But with gravity we have as yet no expression in terms of the same kind of theoretical structures. This is why string theory has been so important. It's the only candidate theory we have that can build a quantum theory of gravity to go together with these other quantum theories.

COOK Report: And the Higgs boson is what is missing?

Newman: The Higgs boson generates mass. It starts with the consideration of how you can have the electro-

magnetic interaction with its massless force carrier the photon, and the weak interaction with two massive carriers - the Z particle that has no electric charge and the W particle that does have an electric charge. So you ask - how you can put together these two forces and have them be two aspects of the same thing?

Peter Higgs, a mathematician, was the first person to write down an expression of the theory where the two interactions are both in the same theoretical framework, and in the theoretical expression he wrote down there is a leftover term which shows that there is this massive particle that is not otherwise accounted for. This particle is now known as "the Higgs". And in the standard way of unifying the electromagnetic and weak interaction which we now know works to great precision there is a missing element which is the Higgs boson.

And even more - when we do precision measurements of electroweak interactions, which are well-described by the standard model which we have done, we can see the effects of the Higgs, even though we haven't seen the particle itself yet. We have diagrams that show the effects of what must be the Higgs and limit its possible mass range. We *almost* know (if the standard model in its simplest form is indeed correct, that is) that the Higgs is there. We just have to go find it. And if it is not there, there are alternatives to the standard model that we can explore.

A Synergy Between High Energy Physics and Cosmology

The slide below expresses a tremendous growing synergy in both physics and cosmology. In particle physics we have come to have a very precise knowledge of the ba-

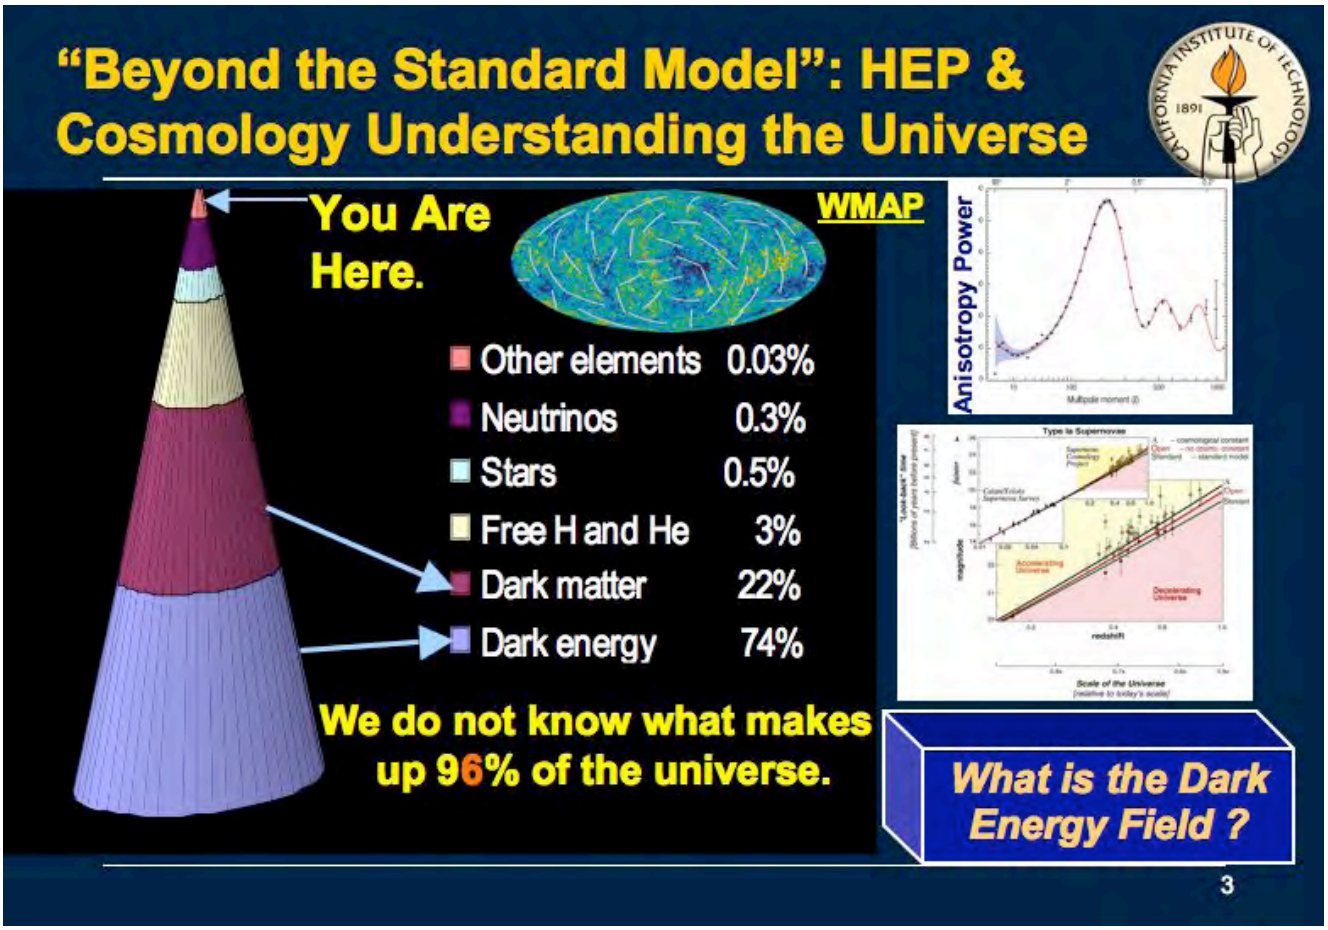

sic particles and their constituents. And in cosmology we have come to a very precise knowledge of what makes up our universe. The flattened globe is actually a picture of the universe and of properties of the microwave background left over after the big bang. By looking at the structure of the microwave background you can figure out a lot about what makes up our universe. Including the fact that it is dominated by dark energy. You can look at the clustering of galaxies on different scales. You can look at supernovae and how they are accelerating away from us. And this has led to the knowledge that there is a predominance of dark energy in the universe and also to pretty much the proportion of the total that is represented by dark energy.

We can see dark matter in various ways by looking at the astrophysical motion of stars or entire galaxies. We know that there is lots of dark matter out there and that normal matter is only a few per cent of the universe. And there is then this exotic and rather well-known picture of the constituents of the universe. The synergy between cosmology and high energy physics comes through the fact that, as the "Big Bang" slide below shows, that the early universe had a very high temperature, corresponding to a very high average energy of the particles the temperature corresponding to a very high energy of the particles exist-

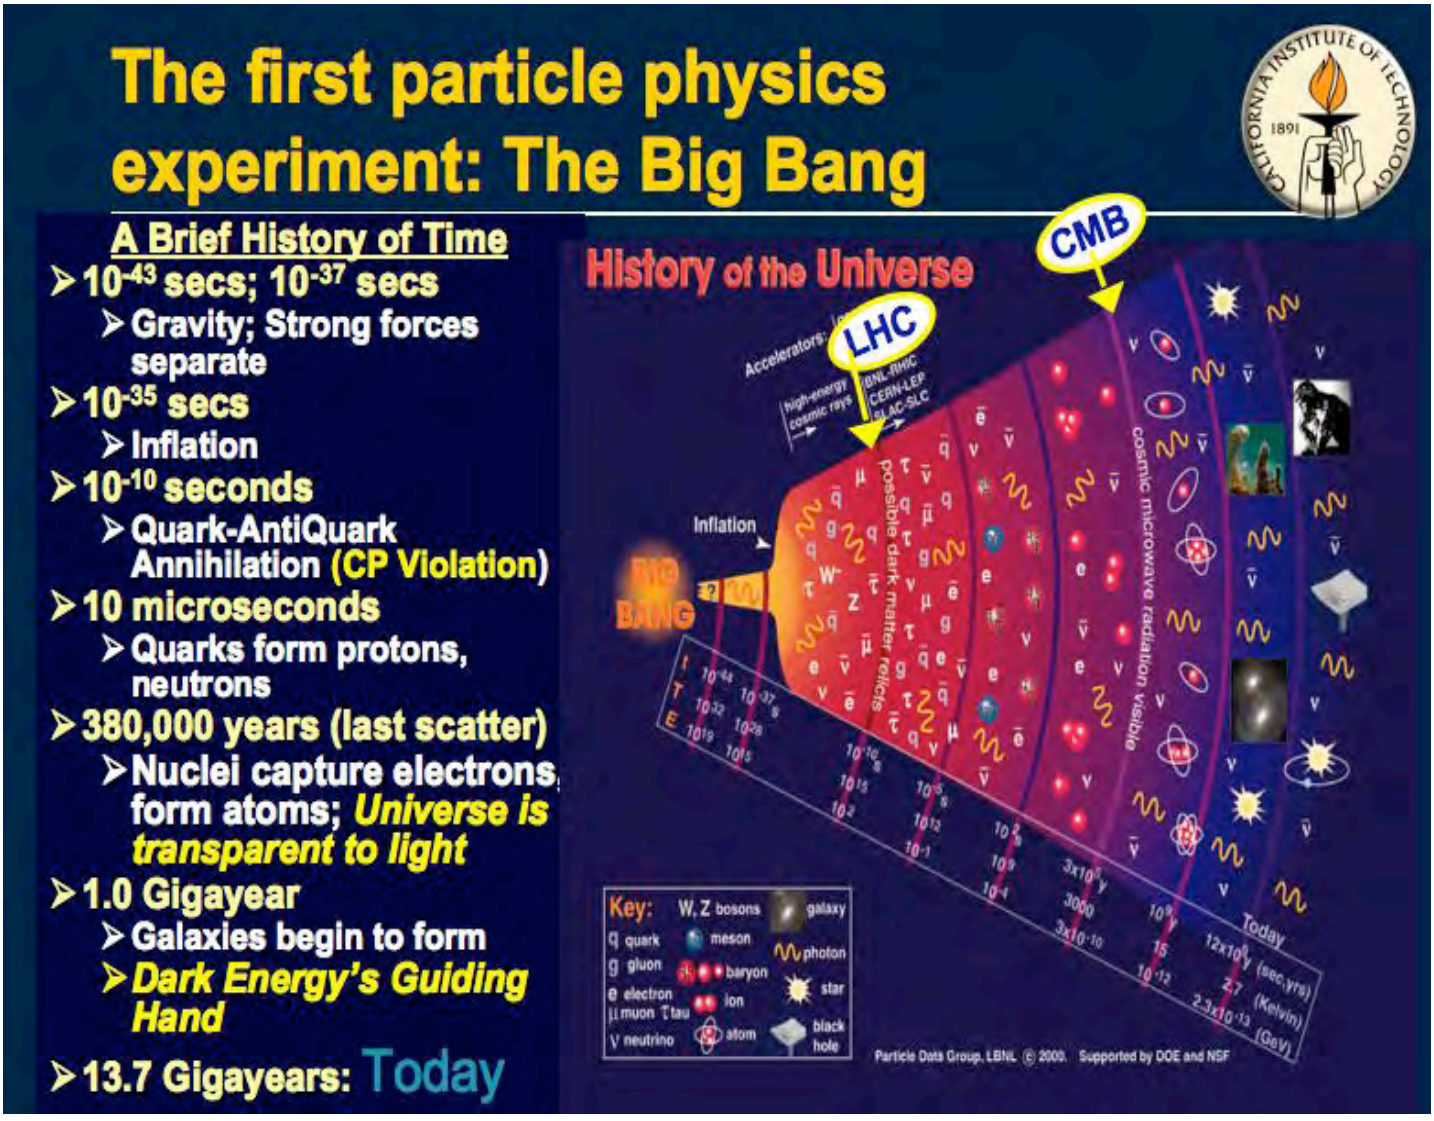

ing at that time. The farther back we go – back to the first seconds or to the minute fraction of a second after the big bang, the higher the energy becomes.

Some of the particles that we can produce with the LHC were last produced by nature about 10 to the -12 seconds after the Big Bang. Therefore, what we can measure in particle physics gives us insight regarding the evolution of the early universe. Particle physicists are very interested in, and extremely excited by the emerging picture of the early universe which in the first instant after Big Bang was such that all the interactions were one; since the energy scale was such that all the interactions merged into one.

If you follow the trend with energy, and hence with the time after the big bang, of the strength of three of the interactions, the strong, the electromagnetic and the weak, as predicted by the Standard Model, you can see that they do tend towards a single interaction strength and hence to one interaction – nearly but not totally.

Then on the slide at the top of this page I have some of the questions that you see in a slightly different form in the Quantum Universe report http://www.interactions.org/cms/?pid=1012346 We know a lot about the universe but

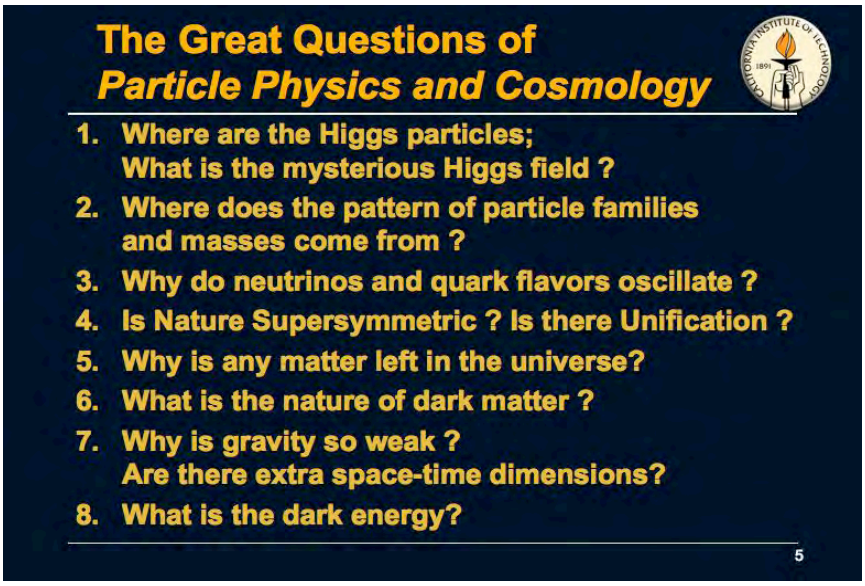

we have so many open questions. After all the whole cycle of modern physics is barely a century old. We have these **Great Questions**. First is finding the Higgs and we think we have narrowed things down to know where the Higgs will be, we just need to find it. There are problems with the standard model. It does not really unify in the early universe. But there are generalizations of the standard model like supersymmetry which do unify, and which avoid some of the other theoretical problems that the standard model has. Supersymmetry is also needed by string theory, and string theory is the one candidate theory we have that unifies all four fundamental interactions – including gravity.

The LHC with its much higher center of mass energy opens up a new field for exploration in order to answer some of these questions.

COOK Report: Suppose we solve these basic problems of high-energy physics and cosmology, will those solutions create opportunities in more everyday areas of physics and nanotechnology in material science and that sort of thing?

Newman: The energy scales here are too high for normal materials. So how you would control these forces to go to some form of materials science? This is something which we don't normally think about. But in the long term, perhaps as long as 50-200 years, understanding these fundamental concepts of the forces of nature and matter may indeed lead to applications beyond our present understanding.

COOK Report: I guess what I'm trying to say is where do

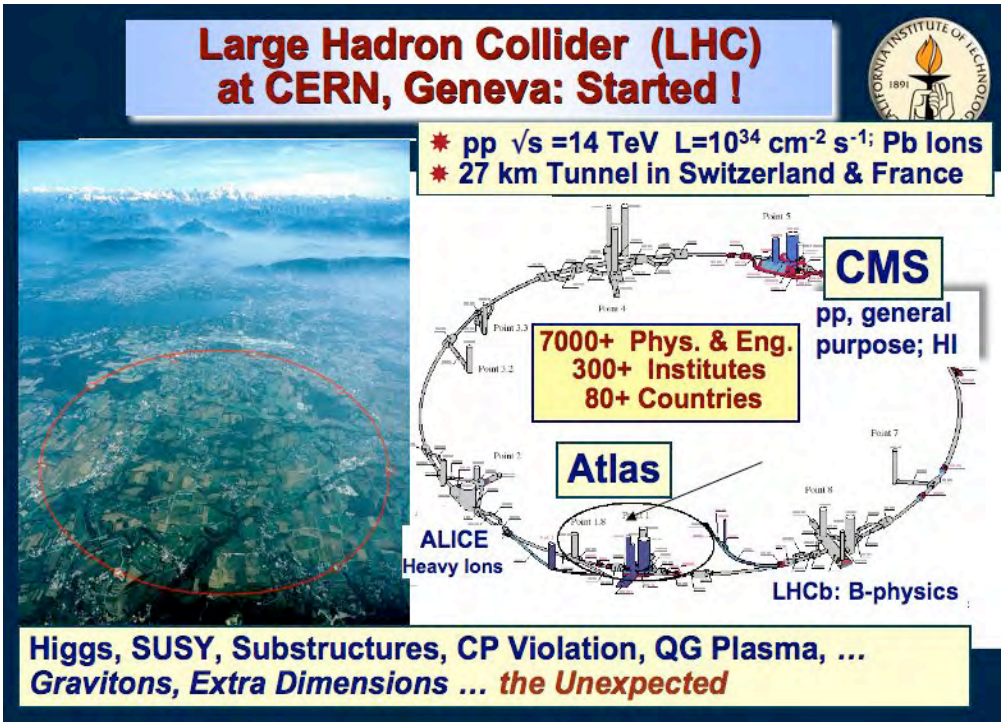

we've been talking about. So there is a continuous flow of short term benefits to society; but the *real* reason that we do this this kind of fundamental inquiry is that is it is an inherent part of civilization. We do it to *know.* It is what humans are all about.

Does it bring benefits to civilization? It does more than that, it shapes the future of our civilization as a whole. For example if there were no physicists to understand electricity and magnetism, there would be no understanding of what to do with electric current, there would be no electronics there would be no computers, … there would be no modern life as we know it.

Can we predict *when* there will be another set of technologies that may create a new form of civilization of which we cannot conceive? I know we cannot. But we also know that one day such a transformation *will* occur. Period.

Then are some pretty exotic things out there, where our ideas have evolved so much recently, such as the nature of space and time. There may also be things that we have not even formulated. But it is possible that what is dis-

you think this research may eventually take us?

Newman: Well we have to have a long view. For example understanding of electromagnetism really began only with Maxwell's equations 150 years ago. It is exactly these explorations that do shape the future of civilization and what it produces. We don't know whether it will be 50 or 100 years until there are technologies that can begin to control forces and particles if not bulk matter on this energy scale. But we think, based on analogies with the past, that that is a possible future. Given that modern physics was only created less than 100 years ago, we can hardly imagine what might happen in terms of further fundamental developments as well as applications several hundred years from now. But if we don't continue this exploration – if we do not do this, then our future knowledge *and* our state of development, will be trivial compared to a civilization (our own or another one that we may one day meet) that does continue to explore fundamental science.

COOK Report: And if you look at the galaxies out there there probably are other civilizations.

Newman: I would think so. The short-term return of carrying out the exploration of fundamental physics is that we do engage with technologies and develop all kinds of systems whether they be radiation hardened electronics or all the other information technology systems that

covered may be unrelated to any of these other conjectures and it is fundamental to the nature of experimental physics that you realize that. Also when you get into these new energy ranges you may see something that you never expected.

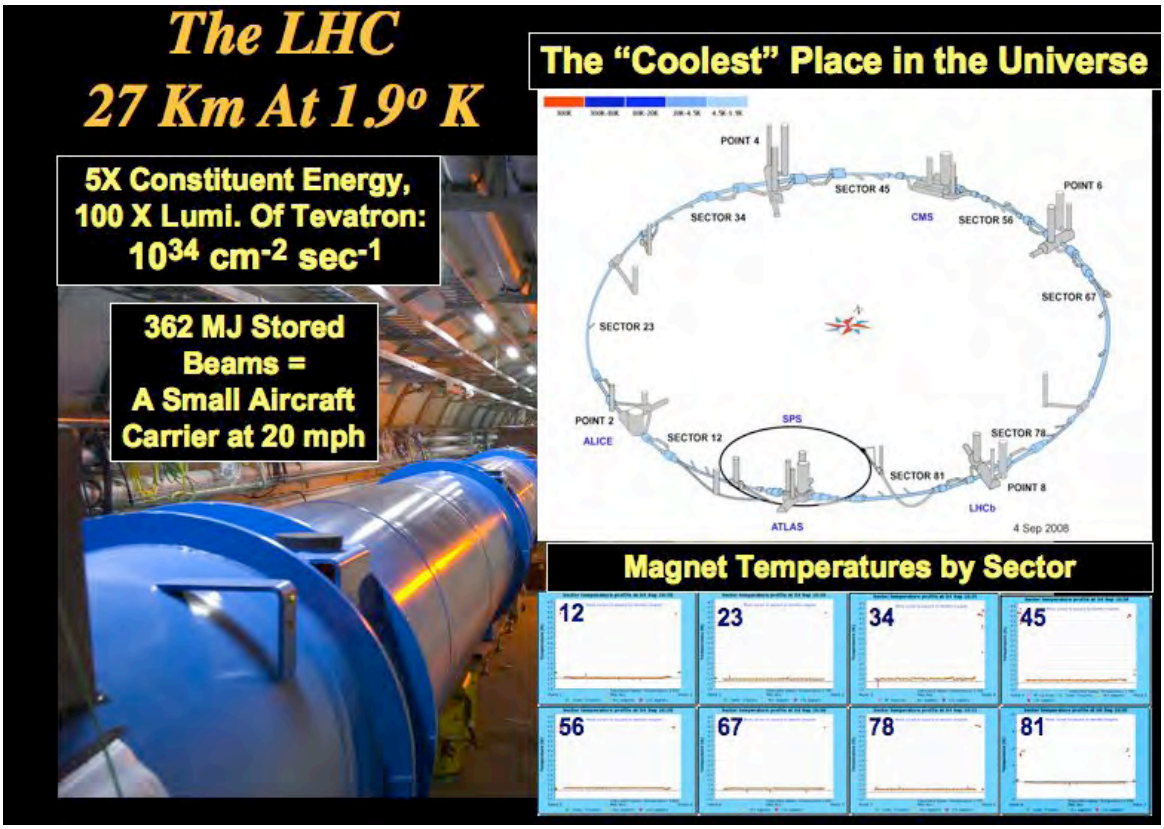

COOK Report: I am looking at the slide to the top right and wondering about the temperature remarks that it makes. Is it saying that the magnets keep the temperature of the space where the beam travels at 1.9° Kelvin?

Newman: Yes. The reason this is the coolest place in the universe is that the average temperature of the microwave background after the Big Bang is 2.7° Kelvin.

Now the next slide that really continues the physics discussion is 14 - bottom right. This just relates to the center of mass and energy of the constituents.

We are not really colliding protons as such. At these energies the constituents of the proton are only loosely bound in the proton compared to the collision energies. Consequently, at these energies most of the interac-

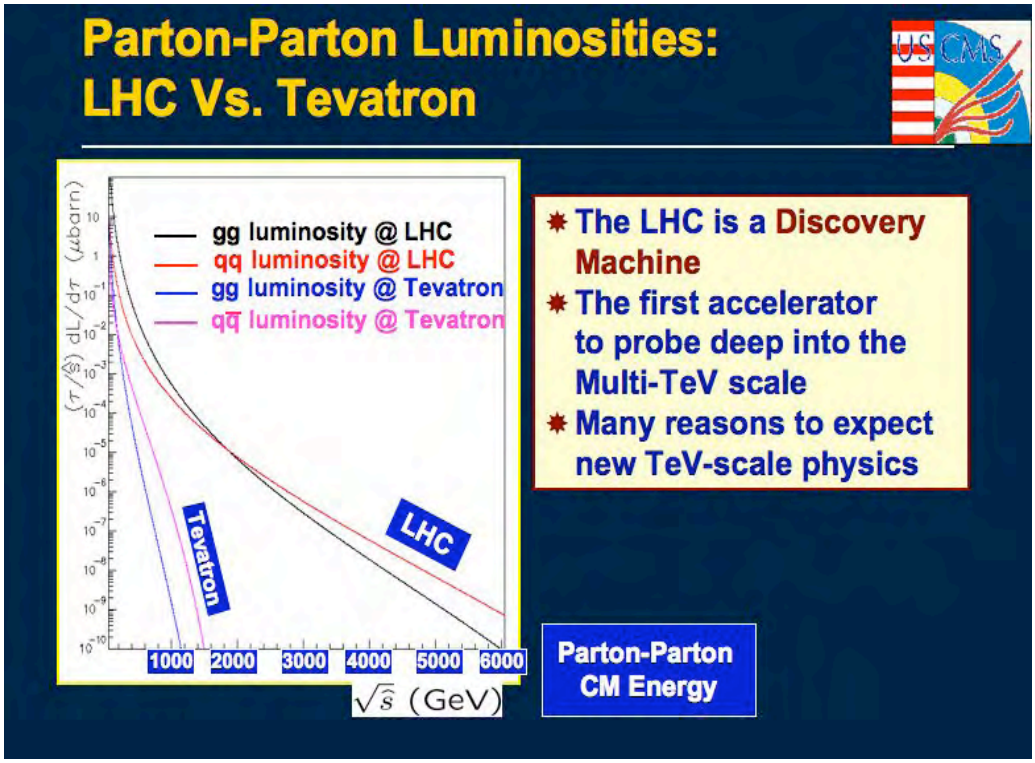

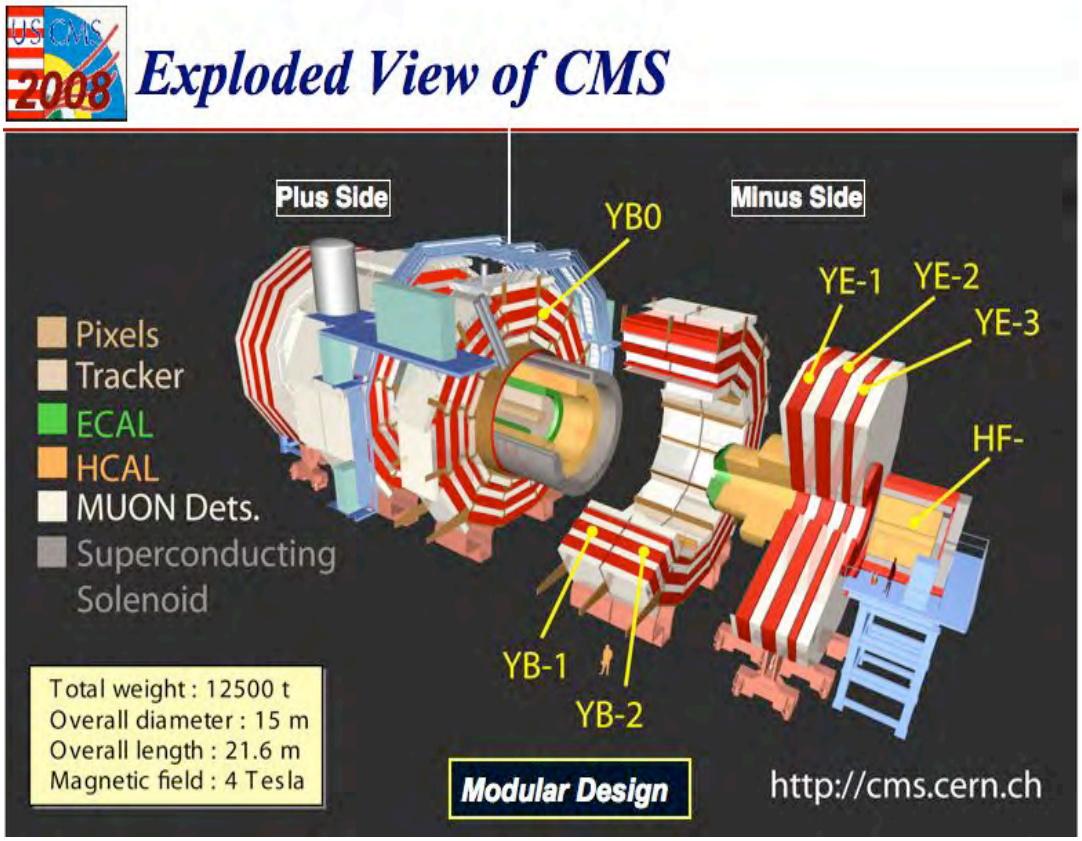

tions are with the gluons that bind together the quarks. You can see the gluon-gluon and quark-quark center of mass energy spectrum we will produce in this. The distributions for the Tevatron which is now operating at Fermilab also are shown, as the two sharply falling curves to the left of the slide. This shows that while the accessible center of mass energy range at Fermilab cuts off between 1000 and 1500 giga-electron volts (GeV) in energy, the LHC's reach extends out to several thousand GeV; about a factor of 5 higher. Consequently, with the LHC, there is an entire new energy range to be explored.

The Collider's Detection Systems

COOK Report: The pictures beginning on slide 10 are fascinating. Would you say a few words about them?

Newman: Slide ten above is the CMS detector which is made up of a number of different subsystems and layers which surround the interaction point in the middle. You are looking at it with the elements cut away so one can see more clearly the parts that make up the subsystems.

When a particle is produced at the interaction point it first goes out through silicon detectors, the first layer of which is called pixels in the legend at the left edge of the slide. Pixels are small rectangular patches of silicon of which there are 66 million in a series of modules. They provide two-dimensional coordinates showing where each charged particle passes through each patch. Next the silicon tracker is outer part of the silicon system. These are strips that, depending on how they are oriented, tell you one coordinate each. With enough layers you can obtain enough information to construct a three-dimensional trajectory of each charged particle. By looking at the bending of the trajectory of each charged particle and knowing the magnetic field through which it passed, you can derive the momentum of the particle.

The next layer in green is called to the electromagnetic calorimeter or ECAL. It completely absorbs electrons and photons. There are crystals that emit light proportional to the amount of energy deposited in them. That light is converted to electronic pulses by photodetectors on the crystals, and the pulses are subsequently digitized. Consequently when any electron or photon goes through these

heavy crystals, there is a whole cascade of these electrons and photons leading to an amount of light proportional to the incident energy. The ECAL then provides very precise measurements of the electron and photon energies.

Then outside that is the HCAL - the Hadron Calorimeter. Particles that go through the ECAL such as pions, kaons, and protons, because they deposit only part of their energy there (unlike electrons and photons), put the rest of their energy in the Hadron Calorimeter which is made up of brass plates with plastic in between that sample the energy as it is absorbed. The amount of energy deposited in the plastic leads again to light, that is then converted and carried back into electronic pulses. That sampled energy gives you a measure of the energy in the HCAL as a whole, which together with the energy in the ECAL measures the incident energy of the incident pion, kaon or proton.

Muons are like electrons but heavier, and for that reason they deposit only some of their energy in the ECAL and HCAL layers. If the muons have enough momentum they get to the outside and there they pass through a series of detectors ("muon chambers") interleaved with the iron layers of the CMS magnet, that measure their coordinates. The white layers on the outside parts of the detector are where we measure the muons.

The gray layer is the superconducting solenoid that generates in the magnetic field. The strength of the field is 4 Tesla, or 40,000 gauss. That is enough to bend particles and measure the momentum even of Tera-electron volt (TeV) particles with 10% precision on the momentum.

The slide below shows a picture of the heaviest, central

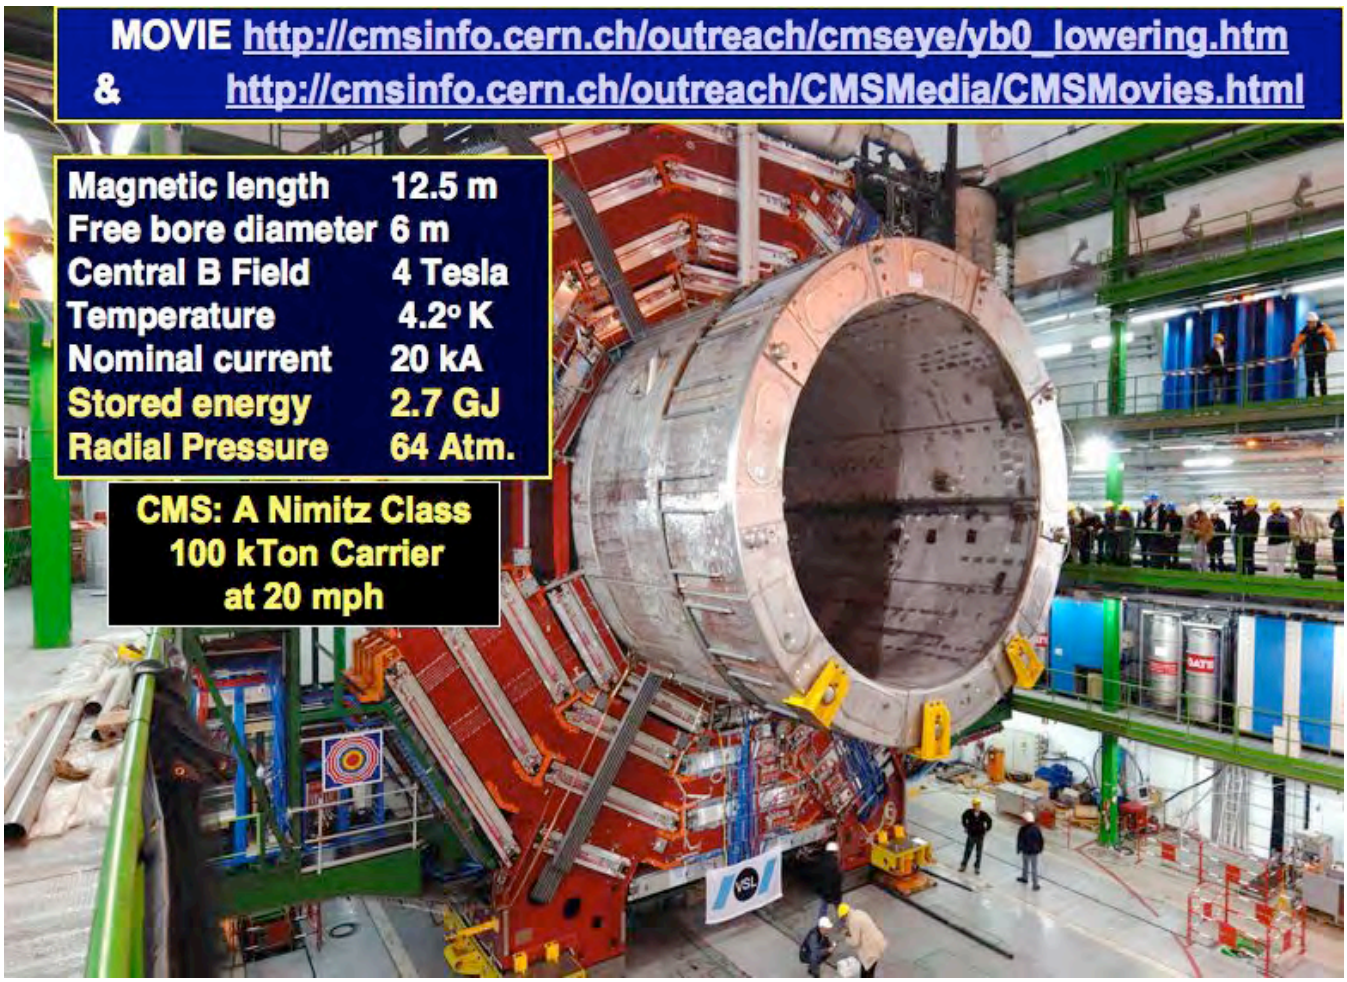

element of the CMS magnet, weighing nearly 2000 tons, and the superconducting magnet coil. You can see its huge scale being more than five stories high. At the top of the slide are two URLs that will show movies of the installation of the CMS. The movies take you quickly through the phases of construction. You can see the iron layers of the magnet which return the magnetic flux, with slots into which muon chambers have been inserted.

The picture was taken in February 2007, just after the pre-assembled magnet element was lowered to the experimental floor. The lowering operation, using a 2500 ton crane spanning the assembly hall at the surface, took eleven hours.

COOK Report: What happened to cause the machine to be shut down not long after it was started? I understood there was a flaw in a weld.

Newman: Yes there was a faulty splice in one of the superconducting bus-bars feeding current to the magnets. A systematic diagnosis of the problem has been done and the affected magnets have been replaced. Improved quench detection instrumentation has been developed and installed along with much greater pressure relief that will prevent an incident of this type from reoccurring.

Let's continue by describing the process by which CMS was assembled and commissioned. CMS is a modular detector with very big elements. The entire magnet was tested for the first time on the surface in July of 2006. It then took a bit more

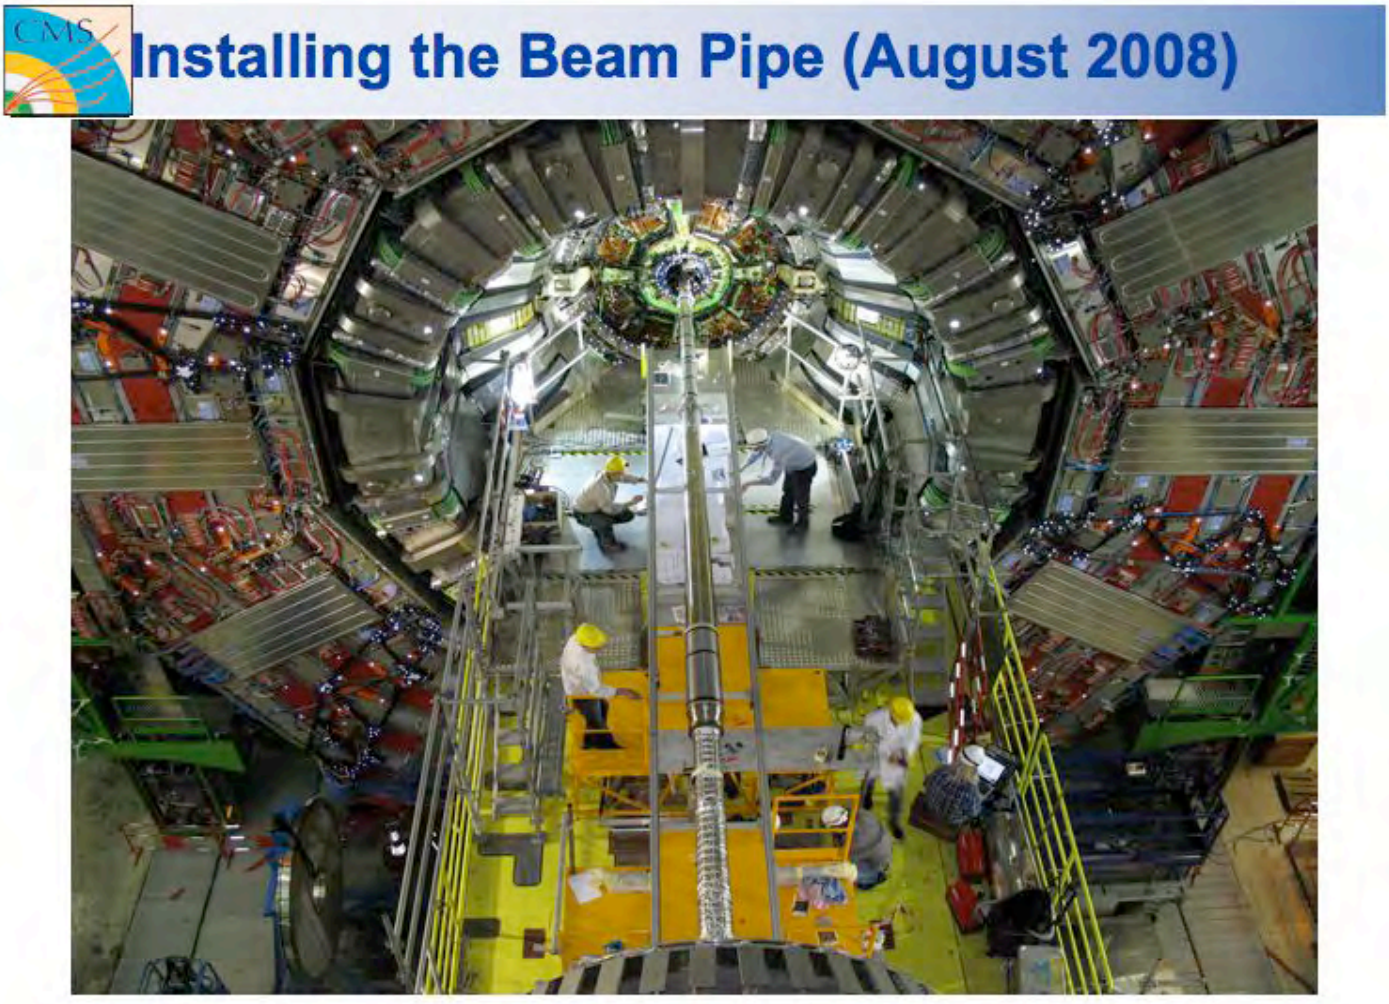

than a year to lower and re-commission all the major elements – all in time for the startup.

Then after lowering all the major elements below ground to the floor of the experimental hall, the next step was to install the beam pipe, as shown in the in the slide at the bottom of the preceding page.

In the next slide below you can see some of the last steps before it was all closed up underground for the first time. You can see one of the end systems which includes the end parts of the calorimeters, like a big nose, being closed up. Actually it was a very tight fit.

The head of technical coordination who was also by training a high-energy physicists said it was like going to IKEA and thinking you only need a screwdriver to assemble what you bought, but finding at the end you need a saw as well. Things are very tight and some of the people you see in the picture underground are observing very carefully to see that nothing physically hits anything else.

In the picture above on the left do you see the octagonal aluminum pieces in pairs?

COOK Report: They look like silver-plated circles in the black area between the two rims of magnets?

Newman: Yes. They are stops. When you turn the magnet on, the iron pieces attract each other with forces of hundreds of tons. Actually the magnet squeezes itself so that parts of the endcap move in by as much as 16 mm. These pairs of stops are there to prevent them from moving in even more.

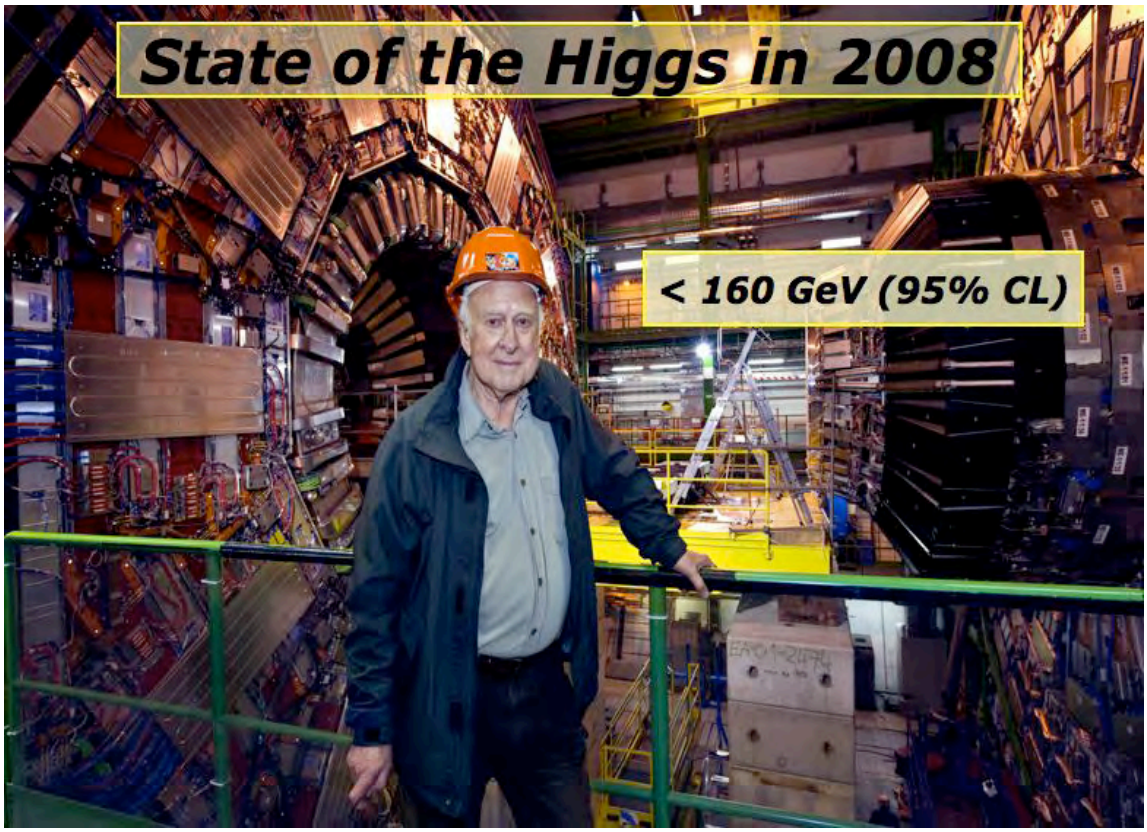

COOK Report: Impressive! This is a device the size of a battleship that is as finely tuned as a Swiss watch.

Newman: Very much so. The tracking detectors are aligned to a few tens of microns when they are installed. But eventually, as the machine is tuned, we will get down to an alignment, using particle tracks, with an accuracy of only a few microns.

State of the Higgs in 2008

Peter Higgs is the mathematician who wrote down the expression for unifying the electromagnetic and weak interactions, that "works" only if there is this massive particle now known as the Higgs boson. That's him standing next to CMS in the underground experimental hall last summer. The next slide is something that my group does in working on detection of the Higgs boson: searching for the Higgs through its decay into two photons. From all the lower energy data that we have, we believe that the Higgs mass is likely to be just out of reach of the previous experiments which were done in the same tunnel in Switzerland but with a different collider. Namely at the LEP electron positron collider that I mentioned earlier.

If you look at all the rates and the distributions in this lower energy data, it tells you that the Higgs mass is almost

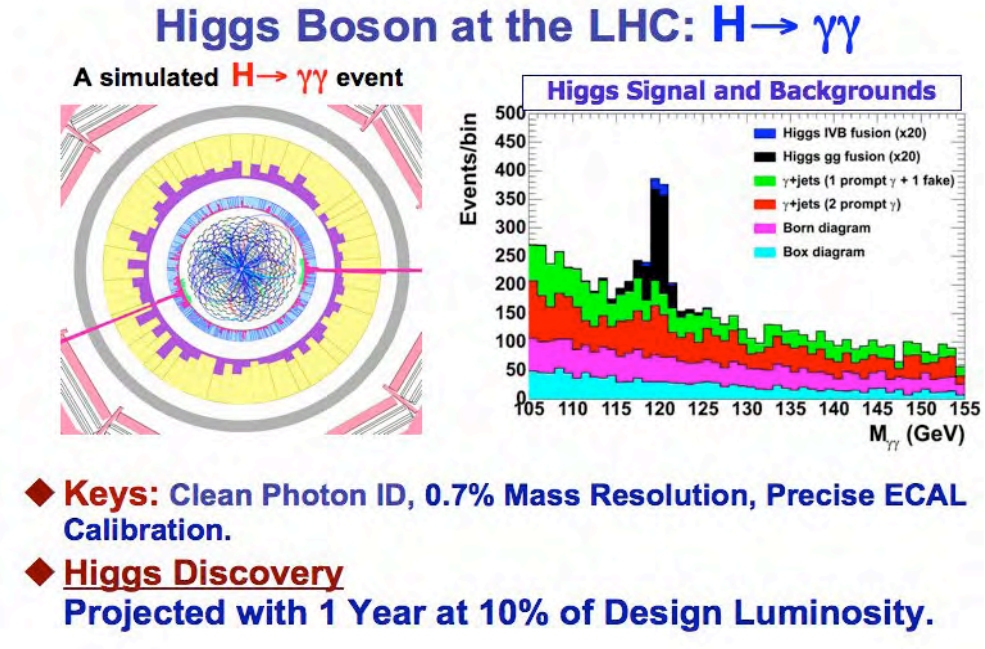

certainly within a range where one of the most likely modes for discovery is its distinctive decay into two high energy photons. It is a rare process but a distinctive one. You can see the picture on the left where you can see lots of tracks on the inside of the CMS detector. Coming out are two high-energy photons which are represented as the peaks in the pink histogram emerging from the electromagnetic calorimeter part.

In each event, you can derive the invariant mass of the photon pair from their energies and angles. By analyzing many events this way, and plotting the distribution of invariant masses, then if the Higgs exists, there will be an increased number of events at a particular mass. In other words a significant excess of events observed at a particular mass will indicate that the Higgs indeed exists.

On the right there is a histogram of the signal, which is in black and in blue since there are two different subprocesses. Below the peak are the background processes each of which produces a broad spectrum of photon energies. Detecting the signal above the background is harder than the picture shows, since the signal peak is magnified by 20 times to make it more visible.

What really happens in the early phases is that as you accumulate data, you find in that there is an excess of events just in this mass range. And also if you look at the details of the photons you find that the background and signal are somewhat, if slightly, different. Therefore you can get a significant signal, and potentially establish the Higgs discovery, before you can actually see a bump in the mass spectrum.

COOK Report: What is meant by the statement: Higgs discovery projected with one year at 10% of design luminosity?

Newman: It will take a couple of years to reach 10% of design luminosity. In the first two years of operation we will have lower luminosity, as the accelerator's performance is progressively increased. At the same time, the ability of our experiment to extract the Higgs signal from the backgrounds depends on the electromagnetic calorimeter's energy resolution, and that depends in turn on its calibration that will improve as we run. The key point about the calorimeter is that it has subpercent energy resolution and therefore sub-percent mass resolution. It is ultimately the resolution, and a calibration technique we developed, that will allow us to separate the signal from the backgrounds.

COOK Report: What would be an analogy by which to understand this? A camera perhaps? A question of getting it in focus?

Newman: A camera is a very good analogy. One can take a snapshot essentially for every bunch crossing online and then decide to set off the flash only a few hundred times a second and actually record those. But when we look at the results, it is like a camera having between 65 and 70 million different pixels. Not only that but each of those pixels has its own different digital information.

Cook Report: Could you explain again what is design luminosity? In non-technical terms.

Newman: Luminosity is proportional to the rate of producing interactions, and that depends in turn on the number of bunches, the number of protons in each bunch, and the size of each bunch. At design luminosity there are 800 million proton proton interactions per second. But in the startup phase of the LHC accelerator there will be fewer bunches, fewer protons per bunch and larger (less-dense) bunches, resulting in far fewer interactions per second which means much lower luminosity.

At design luminosity, the circulating beams have a stored

beam energy of 360 megajoules (MJ) which is an enormous energy to control. The machine group will take time to reach and handle this sort of stored beam safely, and in doing so will reach the design luminosity progressively.

COOK Report: So design luminosity refers to the number of events that the machine is capable of generating when it is fully tuned?

Newman: Yes for each physics process there is a probability of interaction. In particular, what is the chance that a given encounter between a pair of protons will produce the Higgs and have it decay into two photons?

To determine the expected rate of collisions when there is a given number of proton-proton encounters, we use the concept of "cross section". To understand what cross section means, consider that you have two identical black spheres of a certain size and you shoot them at each other. The interaction rate depends on the area of the disk you see when you look at one of them head on; that is the cross section for this process.

For the most common nuclear collisions, the analogous black disk is of order 10^{-12} cm in diameter, which is equivalent to a cross section of order 10^{-24} square cm. 10^{-24} square cm is called a "barn", as in "big as a barn". The Higgs-to-two photons cross section on the other hand is of order 100 femtobarns or 10^{-13} barns, namely ten trillion times smaller. If the Higgs is where we think it is then in a year of operation at full design luminosity you would detect of order 1000 events. And in a year at 10% of design luminosity you would have about 100 events to work with, which is not so much considering that you need to separate these from a much larger number of background events.

Supersymmetry

Let's continue then with supersymmetry. As I mentioned, there are a number of things wrong with the standard model. For one thing, it does not lead to a unification of the forces of nature in the early universe. By 1991, we already knew that if we projected what we had measured at LEP to the energies of the early universe, that we did not get the desired unification, since the projected strengths of the strong, electromagnetic and weak interactions did not intersect at a single point on the graph.

But when the first projections of this sort were done using a generalization of the standard model called supersymmetry, the unification did occur.

In supersymmetry, as illustrated in the slide above, each particle that we know has a partner whose proper-

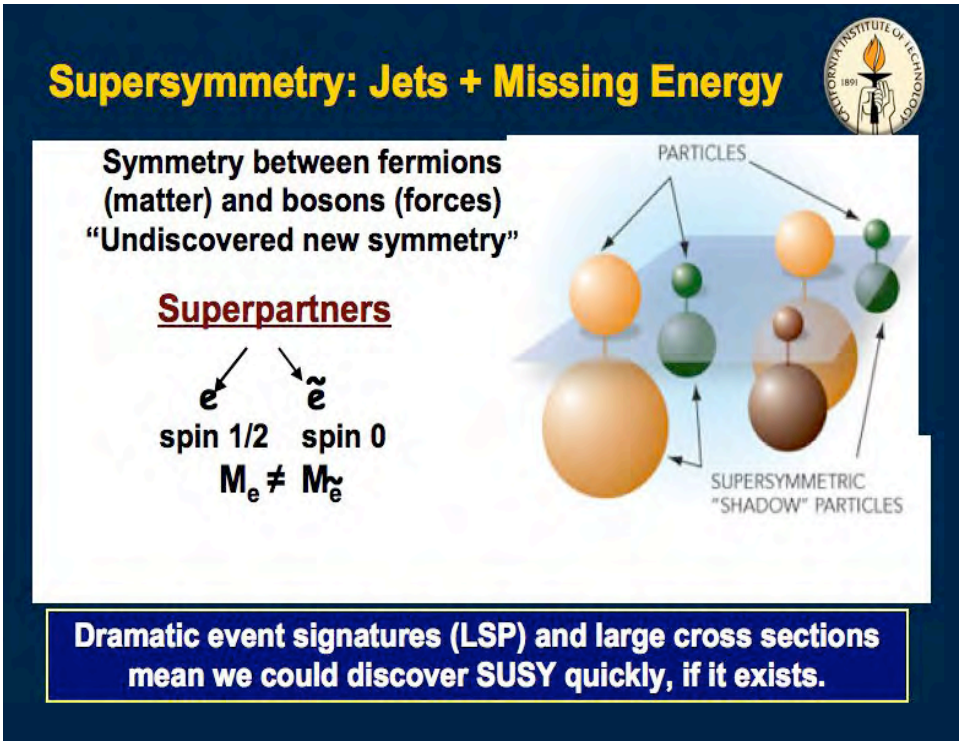

ties would be similar except that their angular momentum or spin is different by half a unit. And so for each electron there is a selectron partner. The electron has a spin of one half, while selectron has a spin of zero.

COOK Report: Do these next slides set up problems that you know you can solve once you have found the Higgs, but problems for which finding the Higgs is a necessary pre-solution?

Newman: We know that some kind of Higgs must exist but it does not have to be the standard model Higgs. The Higgs headings might not decay into Gamma Gamma (two photons). If the Higgs particles fell into the context of supersymmetry, it would not decay as frequently to two photons and we would not see it and we would start to know that something else was up.

So we hypothesize these partner particles but we also must ask why we don't see them and we suggest the reason that we do not is that the symmetry among the known particles and their yet-to-be-discovered partners is only an approximate one. We call this a broken symmetry and hypothesize that all the super partners are heavier and so we have not yet been able to produce them. But if they do exist, at the LHC with its higher energy, we will be able to produce them.

If supersymmetry exists, there will be some spectacular event signatures, as illustrated in some of the slides (you have to use Slide Show to see some of the pictures). One of the first signatures to appear would be events with multiple jets of particles accompanied by an imbalance in the energy in the event. We call this "missing" energy since it means that at least one of the high energy particles produced in the event passed through the detector without being detected. Supersymmetry would have this signature since the lightest supersymmetric particles, which have no electric charge, would interact only weakly with the material in the detector and would therefore not be detected.

Supersymmetry is in one sense a compelling alternative to the standard model. Not only does supersymmetry lead to a unification of the interactions in the early universe, the lightest supersymmetric particle is a good candidate for making up the dark matter in the universe: its weakly interacting nature makes it consistent with the cosmology that we see.

There are a wide range of other new hypothesized scenarios that may occur when the LHC starts producing collisions in a previously unreachable any range. In order not to leave any stone unturned, we intend to investigate many if not all of these scenarios, in parallel.

String Theory

There are a lot of recent developments that relate to how we explain why the unification of the other three interactions with gravity occurred at such an astronomically high energy scale, compared to where electromagnetism and the weak interaction come together. One concept is that we actually live in a larger dimensional space, and one consequence of this is that the unification with gravity will occur at a much lower energy; perhaps at an energy that is within the reach of the LHC. This idea came from string theory. Here researchers found that they could build finite theories but only in 10 dimensions.

One well-developed version of this idea has been developed over the last few years by well-known theoretical physicist Lisa Randall (one of the *TIME* 100 in 2007), working with her colleagues. If this idea turns out to be true, it would lead to some new particles at much higher energy scales than what we have been talking about for the Higgs.

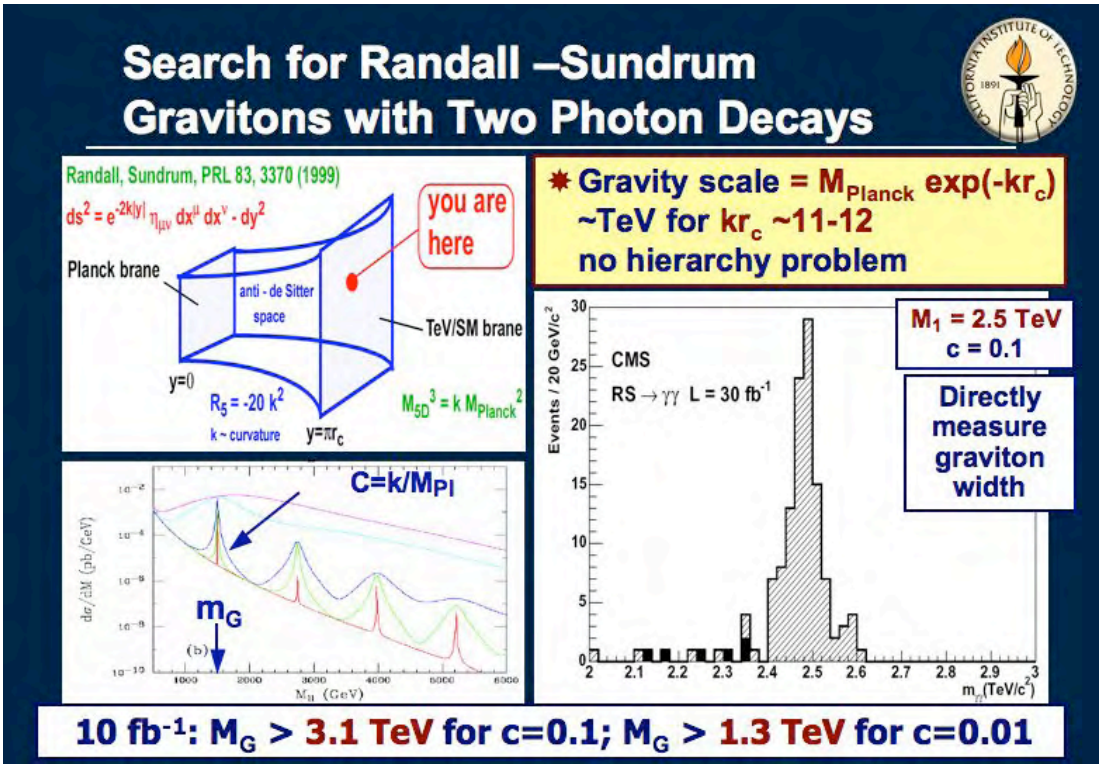

The above slide describes this hypothesis. And the Hawking Radiation Slide below describes in general terms what received all the publicity – namely the search for black holes. Although these might occur in some (relatively unlikely) scenarios, we have spent a lot of time explaining clearly why this would be harmless, even if it did occur.

This is both because, if they exist, they evaporate very quickly and, if they did not evaporate quickly, nature would have created many more of them and, if that were the case, we would not be here in the first place.

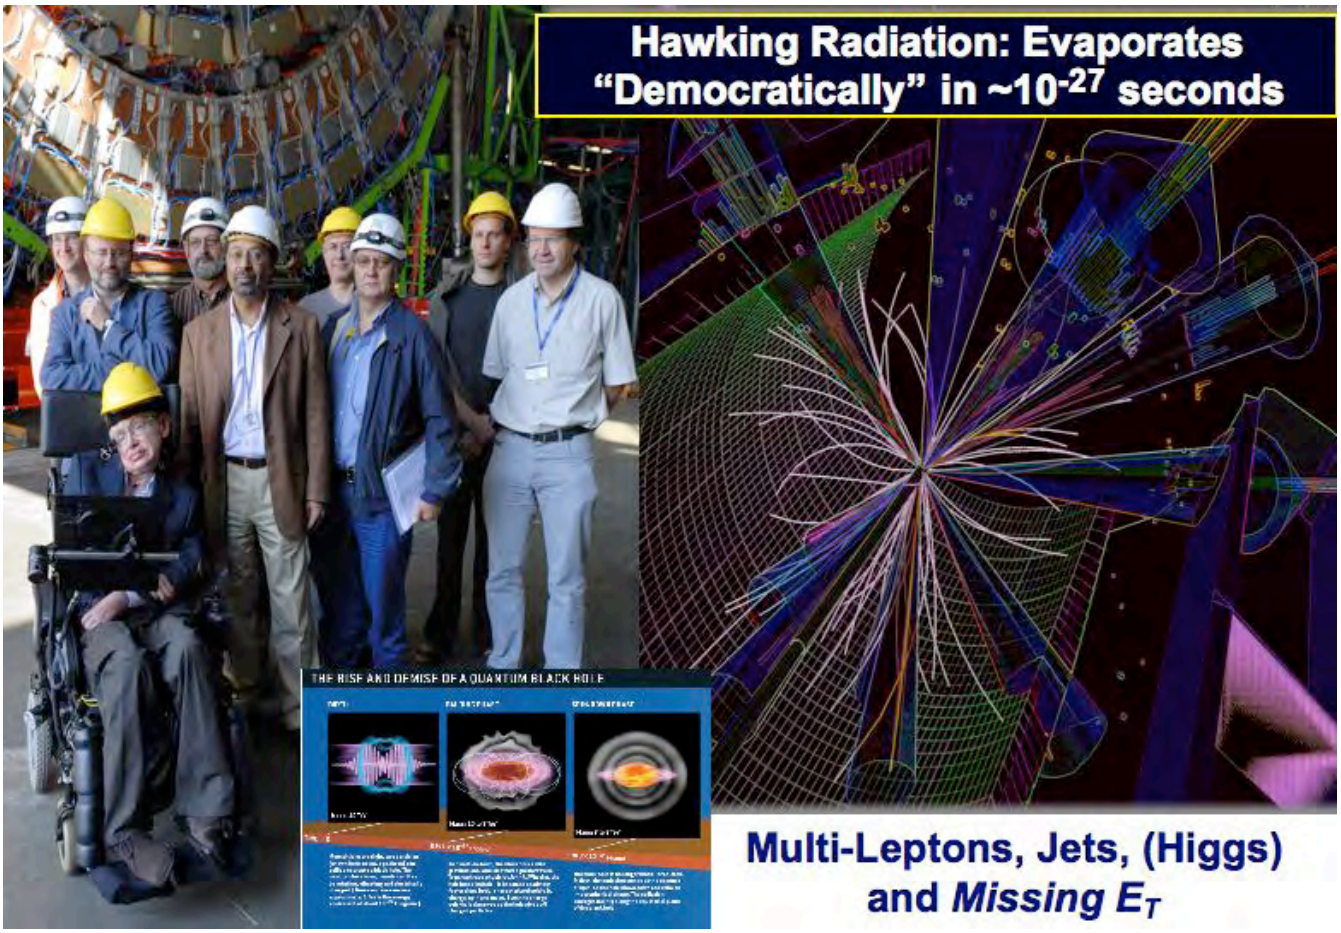

Global Collaboration Mandates Inventive Solutions to Resource Constraint Problems

On slide 20 below we start to get into the links of the LHC science program to the advances in information technology that have been forged because of the need to process, distribute and analyze the massive volumes of LHC data. The slide points to a simulated Higgs event.

Editor's Note: The Preceding slide and paragraph point to part 2 of this interview that will appear in the May issue.

The unprecedented amount of data has demanded a global optical network of unimaginable size, tools to improve the capability of optical channel use and tools for real time audio and visual collaboration among globally dispersed groups of researchers. Part two will cover all of the accomplishments in detail.

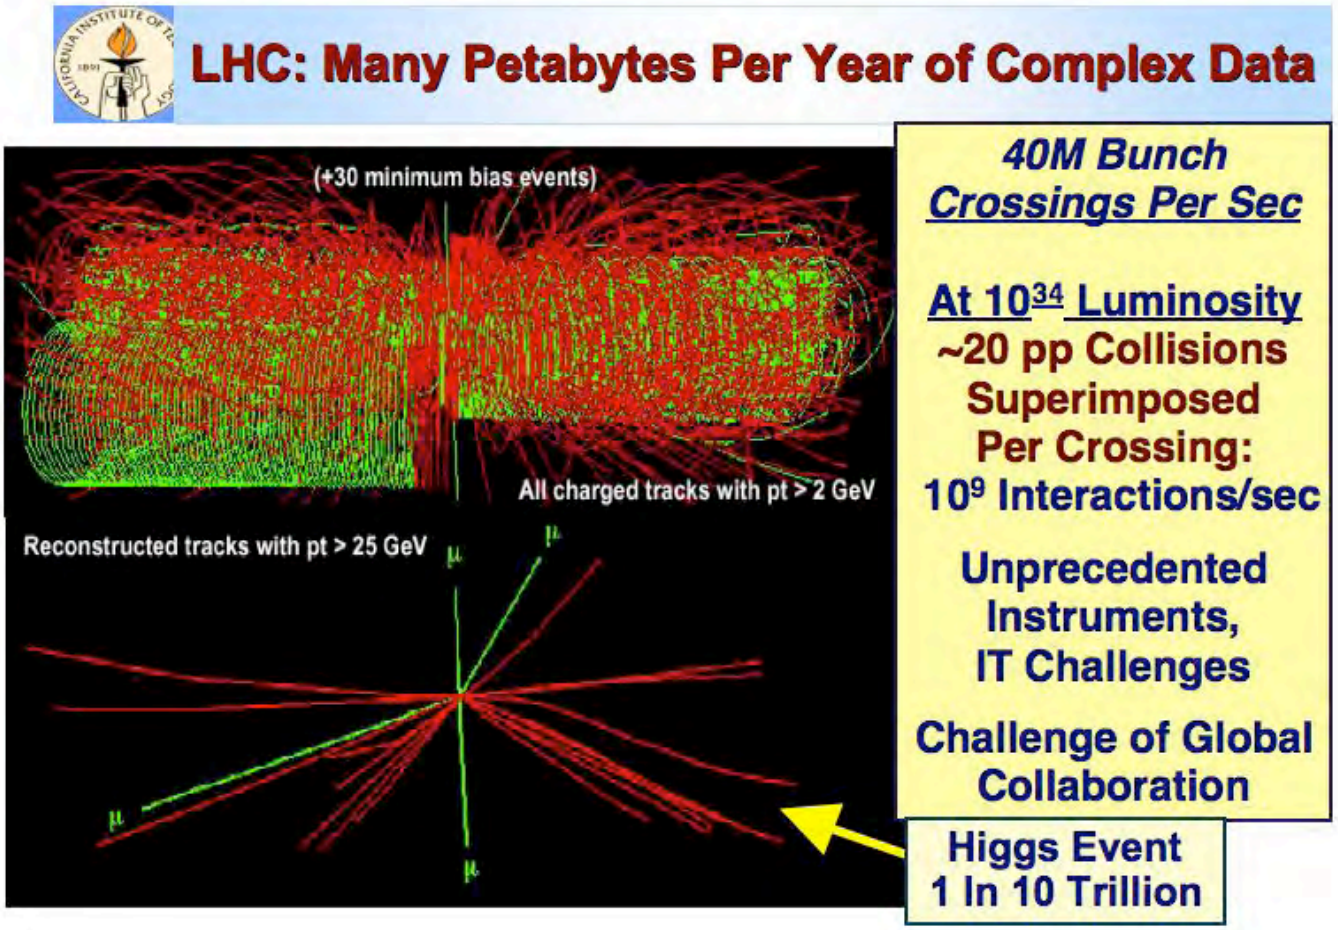

Symposium Discussion January 19 - February 18 2009

What to Do With Fiber Optic Bandwidth? Next Steps for Leadership

Or "Dad, I Want a HyperRoom!!"

From Cook's Collaborative Edge January 22nd, 2009 by Gordon Cook

Jaap van Till is another inventive Dutchman. See his article in the January 2009 ***Cook Report.*** On the Economics of IP Networks list he has come out with a new most interesting proposal.

"Last week I introduced in Holland the idea not of a new device/gadget, but of a cluster of gadgets which I call a "Hyperkamer" in Dutch or "HyperRoom". My intention is that students and teachers get a serious number of tiled flat screens on a wall in their room at home, a VolksWagen-version of the OptiPuter screens and networks with which scientists are experimenting. These are all controlled by App's on a smartphone, by pointing at parts of screens for TV, virtual models, games, documents or manipulating images on the screen wall in total. The smartphone can also be used in-house to control all other gadgets and boxes in the room and all electric functions in the smart home as well. In essence, the ultimate single remote device for the complete house. And everyone in the house has got one, to boot."

"The name comes from the fact that the present Net-Generation (age 13-30) is already HyperConnected, intensively using more than four online communication gadgets. **The essence of my idea is that students in their HyperRoom can interoperate these gadgets and can learn and cooperate with lecturers and team member students in their room and at multiple locations simultaneously. By manipulating the info on the screens together, they can create economic value, synchronize and synthesize their different contributions and visions for projects and mashups to design assignments, products and new solutions.**

I have asked the Netherlands government to fund a project (het HAN Hyperhuis Project) to let my smart multi-talented students at the HAN University in Arnhem near Amsterdam define and design such a multi-user networked virtual creative class environment for energy efficient use in their own rooms at home. I can imagine that Apple and Google might help fund this project too. Maybe this is the metaphorical Car of the Future? It is not hard to imagine that the Oval Office will be a cool and well connected HyperRoom soon, too. A room with a view indeed. Why does this professor do this dreaming ? It is the least we can do to help Steve Jobs stay connected while recuperating and…… I want one myself, don't you?!"

Cook's Edge - makes sense. It syncs with Harvey Newman's comment from the same January issue that "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." See also Interview with Harvey beginning page 1 of this issue.

John Waclawski has been eloquent about the need for communications interoperability of all manner of networks in the home…. and scathing in his comments on standards groups as too often bastions on non interconnectivity. I suspect Jaap is quite correct that hyperconnected kids would relish opportunities to unlock their digital gadgets and make them interconnect. The world needs to be more democratically productive and in an age of commodity hardware, open source software, open interconnectivity, the next step is to broaden the open source nature of the optiputer by showing kids ways to interconnect their own communications devices in their own homes and schools. The desire is there as this picture set from Japan shows. http://www.dannychoo.com/adp/eng/1653/Japan+Optic+Fiber+Internet.html

On Arch-Econ **Jaap** explained further: "the 'wall' is only one component in the room-network I propose. In the case of the HyperRoom just ask any young intelligent person what he/she has in his/her room at home now: Laptop, TV, beamer, game console, cellphone, video recorder, DVD player, CD player, MP3 player, books, coffee machine, loudspeaker, iPod, webcam, photo camera etc, etc. etc.

Do they interwork ? No. Can he /she for instance get the TV images on the laptop or the Internet Youtube or delay-TV images on the TV set?? NO

Obstacles: content owners, formats, proprietary technology. Island design. Just like wide area networks in the pre-internet era.

Students can together design around these obstacles and build first class work/study/cooperation environments (micro-internets) for themselves if they band together and make a fist. And by using technolgical solutions which are available and under construction by the big-science guys (as depicted in the Feb 2009 COOK Report issue). All I suggest is take that technology and put it in the hands of the young. Starting with the college- and university students. My slogans give focus and banner text to such movement. And - at least in the Netherlands - we have the broadband Network infrastructure in place to make this move. So we can stay ahead."

Cook's Edge: So says Jaap. Smart man.

So where is ICTRegie? And an American equivalent program linking Ed Seidel, Harvey Newman, and Irving Wladsawsky Berger.

Anyone who thinks this is just "cute" should look at Tom de Fanti's December 14 2008 presentation of the state of opti-portals and green computing. http://www.is-uc.org/2008/img/DeFanti-ISUC-GreenLightFinala12-14-08.pdf And then finish off with the consumer grade tiled displays already being sold. http://www.digitaltigers.com/displays.asp?r=MaximumNumberScreens&value=8&bhcp=1

More About Risks of Emerging IPv4 Markets

On January 27 **Harvey Newman**: Various sources, such as the TERENA Compendium, show that the uptake of IPv6 is extremely slow. Yet IPv4 address space continues to approach exhaustion.

Here is a position paper from Network World and NTT explaining IPv6 and the issues.

http://edge.networkworld.com/whitepapers/nww/pdf/C010_EG_IPv6_1208v2.pdf

Goldstein: Yes, it's hilarious. And while the title is "not if, when", the first article is all about how NATs are central to its deployment. Wait a second... the irony is too much, and they can't even paper it over. They have to destroy the village in order to save it, but by gum they'll save it. I do like the Geoff Huston quote, though; his lack of enthusiasm for V6 is obvious.

And their article about federal adoption forgets to mention GOSIP, which would have migrated all federal networks to OSI by 1987 or so, by the same logic. Young reporters...

The article on Page 25 is a real howler. That article is so ridiculous that the mere act of trying to refute it might be degrading. Just the initial 10:1 ROI is so preposterous that it is probably self-defeating. Of course I see similar claims of the value of "broadband" investment, so making up numbers like this out of whole cloth is common practice.

I love how at the top of the article's p4, they cite "IPv6's self-discovery capabilities" -- do they mean the MAC address in the IP address? That's a known security risk as well as an overall Dumb Idea, though it seemed clever in DECnet 25 years ago (the previous DECnet buggered the MAC address to put the 16-bit DECnet into it, so they were still figuring out broadcast topology subnetworks).

And by "peer-to-peer nature", do they mean that IPv4 is really SNA LU2.1, a polled master-slave protocol? Then they site global addressability of zillions of devices as a security feature, when it's a risk. (Do you really want a port scanner in Kyrghizstan to have direct access to your home automation system?) And they claim multicast capability, which just isn't there in any meaningful way.

The v6 crowd has degenerated into self-parody. On the bright side, seeing that NTT paid for this, it's proof that the Japanese aren't all so smart after all.

Newman: I still see that we face IPv4 exhaustion, and given they have insufficient space for historical reasons, China is building an IPv6 network on a large scale (what else can they do?). So if IPv6 failed, what is the alternative?

For example, what about simply having an extended address field, and one bit linking to that field if it exists, as the minimalist solution ? ASICs in (all) routers and switches could easily handle that.

Goldstein: The correct longer-term answer is to sit back, look at what we're *really* trying to accomplish, and start over without the burden of the past. We don't want IP with longer addresses; we want applications to be able to connect. (Stanley Works model: Holes, not drills.)

The short-term answers are twofold. One, existing IPv4 address space should be used more efficiently, and blocks should be transferred from

those who have too many to those who need them. This is what the new transfer policy, which has been discussed here recently, is about. Policy or not, this will happen; market forces are more powerful than ICANN-designated bureaucrats.

Two, NAT should not be viewed as an aberration but as a proper component of the stack. Translation and encapsulation both have a role. Each "network" should thus have its own address space, and the global network should be viewed as a catenet, let's even call it an internet, of these address-domains. The global space should be kept for the backbone networks and for public servers, since the legacy protocols insist on (this is a mistake in the current architecture) doing the IP address lookup in the application, rather than passing the name down the stack. Any application that is not "NAT-friendly" is thus defined as broken. That means no IP addresses in the application layer. (Old sentimentalist that I am, I'll tolerate FTP's mistake within FTP per se, as I used to print my own work through that same PTIP. But nothing newer.)

Harold Feld: Ignorant question: can one run both v4 and v6 simultaneously?

Conrad: Yes. That was, in fact, the IETF approved "transition" plan.

Feld: this would have serious congestion/traffic issues, . . .

Conrad: Not really. If you're able to connect via IPv6, that connection would not be done over IPv4 so there would be no appreciable difference in traffic. There are some routing table growth issues, but that was a hard problem that folks chose to ignore.

Feld: . . . but would encourage a gradual phase in to potentially reach critical mass for v6.

Conrad: Unfortunately, there was (is) no reason to move to IPv6, so there was (is) no point in vendors/service providers spending money to add IPv6 support, so there was (is) no reason to move to IPv6. Now we're running out of unallocated IPv4 addresses and the "dual stack" transition model, which implicitly assumed that IPv4 addresses were available to number the IPv4 side of the dual stack, has some challenges. As a result, the IETF (after having come up with and then deprecated a first attempt) is now again looking at IPv4 to IPv6 translation.

All a bit of a mess, really.

Tom Vest: The answer to your (Harold Feld's) first question is yes, and that "dual stack" capability was supposed to provide a form of backward compatibility during the anticipated transition. Unfortunately for everyone who might want to participate in a transition after this year (+/-), that capability requires both IPv4 and IPv6 address resources, and the former are going to be "challenging" to come by.

The answer to your second question is probably no; congestion will not result -- would not have resulted even if the dual-stack strategy had worked. However, the demand multiplier effect on routing system carry capacity might have posed a problem, eventually -- just as it might, eventually, if IPv6 ever actually takes off - just as it is certain to if service providers accelerate/intensify their current practice of de-aggregating the address blocs they receive and announcing them as multiple, smaller prefixes.

Absent the specific combination of IPv6 plus change of operational practices -- or some other currently unforeseeable alternative technology - or something even uglier (e.g., an overt routing cartel) -- that current practice and the problems that result will continue to escalate, since it's a natural/inescapable by-product of the "more efficient" use of IPv4

that Fred and other address market boosters have been advocating.

Interesting times ahead...

COOK Report: Is there a new aggregation business out there for someone who studies allocation and figures out what blocks have enough aggregatable unused addresses to be routable and then convinces people to transfer title?

Conrad: ISPs are already in the aggregation business. I suspect this will simply become more explicit. Some have expressed concern that speculators will go out and buy up all the allocated-but-unused address space. While this may be true, it sort of misses the point -- the only time this becomes important is ***after*** the IPv4 address space free pool has been exhausted, in which case most folks would already be unable to obtain IPv4 addresses. The fact that speculators will have locked up the allocated-but-unused address space merely maintains the status quo...

Vest: Actually, the status quo is likely to change in one huge way, with multiple far-reaching effects, regardless of whether the market "works" and liberates lots of additional IPv4, or absolutely fails to motivate current "reserve address holders" (i.e., hoarders), or is fatally distorted by speculators, etc.

The next time a lawsuit arises over control or use of IP address resources (it has happened several times in the past), at least one side is going to have recourse to the protections afforded by whatever property law is in force in the relevant jurisdiction. Before transfers, resource users were not considered to "possess" IP number resources in the way that makes the phrase "possession is nine-tenths of the law" such a perennial favorite. Once transfers start happening, and lawyers and accountants are forced to reckon with the implications for both transfer transaction participants and everybody else, this will no longer be true.

The first likely impact will be to render moot any policy-based requirements or restrictions on IP address transfers themselves. The really big impact will come when this starts to undermine participation in (any) shared public registration database -- which to date is the only thing that has preserved the presumption of uniqueness that puts the "public" in "public IP addresses". If that goes, the only things that might be able to put it back will be property law and/or public regulation -- which will have to be reconciled with other laws about privacy, cross-border trade and direct investment, etc.

Both alternative mechanisms are almost universally nationally defined, so the results -- and any guarantee of uniqueness itself -- are quite likely to vary substantially from country to country. Voila, instant national segmentation of the Internet.

Some market advocates publicly reject this scenario as speculative, but a fair number (also) embrace it as the intended outcome.

John Levine: It's not implausible that this could push v4 exhaustion back by many, many years. The numbers I've seen suggest that something like half of the currently allocated space isn't actually in use. I know that I'm using less than half of the /24 my ISP allocated me, but as it stands there's no incentive for me to do anything about it. If they offered me a modest inducement to renumber into a /25 or /26 I'd probably do it.

List member: Guys help me to understand a little more clearly if you have a 24 and only need half, you sell it to big isp.....err transfer it to big isp..... but that big ISP

won't dare try to route it globally until it has aggregated enough addresses from the surrounding block to get a 21 or a 19.... something large enough to be globally routable..... yes?

Levine: Not really, it's likely to evidence itself as a lot of intranetwork shuffling around that doesn't affect the global route tables.

I have 208.31.42/24, which is actually part of a /11 assigned to Sprint, from which they assigned a /21 to my ISP, from which they assigned me the /24.

At some point my ISP is going to go either to Sprint or to ARIN to ask for more space and the answer will be that it will cost them a zillion bucks. So if they can persuade me to shrink down to, say, a /26 and renumber a bunch of their other customers into the rest of my chunk, that could free up space a lot cheaper than getting new allocations.

Vest to **List member**: The likelihood of actually observing that in the wild approaches zero. For the scenario you describe to be possible, the /25 that you want to sell would have to be directly adjacent to the one lone prefix (of no more than /25 length) that is directly adjacent in the numeric sequence of IP addresses. If it's not adjacent, it can never be aggregated. If it's adjacent but larger than /25, it could already be routed today, and tacking on your idle /25 would neither be useful nor, likely, possible.

Vest to Levine: Unless you have an ironclad contract that explicitly guarantees a fixed price for your current /24 in perpetuity, I would assume that that "persuasion" is most likely to take the form of a substantial increase in the cost of your public IPv4 the next time you have to renew.

Levine: Oh, they wouldn't do that. They're not that kind of phone company. (Fred can tell you what kind they are.) When I shut down my T1 as I moved out of the country, I called them up and asked if they'd just reroute my /24 to the ISP down the road where I put my server, who also happens to be their customer. "Oh, sure."

Vest: Sounds like a great company. However, it also sounds like a company with increasingly valuable but "underperforming" assets -- i.e., the kind that would likely be regarded as a very attractive acquisition target by a larger, more aggressive ISP, or perhaps a speculator.

Earlier Vest: Your ISP may or may not face the same kind of persuasion from Sprint.

Levine: They've grown over the years so they actually have more space directly from ARIN than from Sprint, and they have their own ASN and are multi-homed. I dunno what they'd do if Sprint put the screws on, they probably don't know either. As likely as not they'd politely ask if I'd renumber to help them avoid a big price increase, and I'd politely say yes.

Vest: Fair enough.

Levine: Sprint, who presumably also has IP space needs, could offer to pay my ISP for any space they give back, with a modest payment per /24, increasing more than linearly for bigger chunks.

Earlier Vest: The resource transfer proposals currently implemented or under consideration by the RIR communities provide no new options (no need to call them "rights") for IP address users that do not have a direct relationship with the RIRs themselves.

Levine: Not directly, but once there's a market for inter-provider transfer of IP space, I'd expect internal transfers to happen at similar prices. Wouldn't you?

Vest: I expect that no one currently holding transferable IPv4 is going to let any of it go at prices that would ap-

peal to anyone without a critical need -- i.e., very very high prices. If you are planning to offer high value commercial IP services, and therefore capable of justifying a commercial-scale capital investment in inputs (I am assuming 5-6 figures, minimum, for a true provider-independent /24 -- but at this point it's anybody's guess), then that would make you like any other "new entrant" -- just one that hopes to get address space from your former dedicated access provider.

Earlier Vest: If, as you say, your ISP received its address space from Sprint, then your ISP is a price-taker for IPv4. Barring any explicit contractual prohibitions, if Sprint were to unilaterally elect to raise the cost of your ISP's address space, or even to reassign most or all of it outright, the only recourse that your ISP would have would be to cease being a Sprint customer for IP transit (or whatever paid service the /21 is bundled with) and try to find a better deal with another upstream provider.

Levine: Depending on how efficiently they allocated space it could go either way. (I am, I believe, my ISPs oldest customer with my original /24 allocated in 1995, and renumbered once when they switched upstreams. They allocate a lot more efficiently now.)

There also appears to be great chunks of little used space from the early days ss Xerox has 13/8, HP has 15/8 and 16/8, Apple has 17/8, Ford has 19/8, Computer Sciences Corp has 20/8, Halliburton has 34/8, Eli Lilly has 40/8, and so forth. I assume they're each using some part of their space, but I expect a lot of them could renumber into a /14 or less if they knew they could sell the rest for millions of dollars.

Heck, even MIT has a 18/8 and U. of Michigan has 35/8 (and I bet Stanford feels really stupid for being good doobies and renumbering into a /14 in the 1990s.) There's plenty of space in large chunks if people wanted to make it available.

Vest: I agree with you that if a transfer market persists, it's quite likely that some of this presumably underutilized address space will find its way into the marketplace.

However, as long as IPv4 remains an absolutely essential, absolutely non-substitutable requirement for attaching to the Internet, there's no reason to think that the current holders are going to liquidate their newfound wealth in any way that doesn't maximize advantage to themselves. However, even if they do choose to act in an "irrationally charitable" ways, for example, by making their surplus IPv4 available to the market at "cost" (i.e., between $zero and the labor cost of renumbering), then the first takers are likely to be speculators, who will be only too happy to flip the address space and bring it back into line with whatever the market will bear. They don't call them "market makers" for nothing ;-)

Levine: Nobody said it would be cheap. We can only be sure it'll be cheaper than rewiring everything for IPv6.

Vest: Really? How much will you be willing to pay for IPv4 personally in order to make IPv6 unnecessary? How much should everyone be willing to pay? I suspect that there is some IPv4 price point at which you might change your mind. However, even if there is it won't matter -- because the decision to make IPv4 substitutable, via IPv6 or something else, will be made by the people who set your IPv4 prices, and the prices for your IPv4-based services. So long as some people are willing to pay for IPv4, and everyone else has no choice but to settle for a degradation of services (i.e., ever-increasing use of NAT and RFC1918 space) to accommodate those who are willing to pay the IPv4 premium, the Internet will become both

progressively more expensive and progressively less useful for everyone, except perhaps the inheritors of RIR-era IPv4.

I'm personally hoping for "magic" to prevent this from coming to pass, but I'm not very good at blind faith :-\

Earlier Vest: In all likelihood each [current holder] has already been approached numerous times by aspiring brokers/middlemen offering to help them capitalize on this once-in-a-lifetime opportunity.

The first time I heard about such an inquiry was over a year ago.

But Aren't the RIRs Restricting Transfers to 24s?

Jan 29 Rudolf van der Berg: Isn't that also up to the RIR's. If the policy is that the only blocks they allow to be transferred are /24's then that is the bottom limit?

Vest: Your question is a bit ambiguous, so I will try to answer for all permutations.

What constitutes the longest (smallest) "generally routable" prefix is an operational matter, determined by the commercial practices of individual routing services providers, completely outside of the RIR process.

The fact that this informally defined value has been quite different (i.e., much smaller) than the collectively defined, RIR policy-mandated minimum IPv4 allocation size has been a frequent source of criticism by market advocates, who claim that the RIR system has failed in its second core mission, i.e., to check the rate of inflation in demands placed on finite routing system capacity. The flaw in the critics' logic is that this gap between policy and practice proves that the policy was completely ineffective, had no value at all. This is sort of akin to the leap that people often drive 10-20 m/h km/h over the official speed limit on highways, ergo it would make no difference -- average speeds would remain exactly the same -- if all speed limits everywhere were abolished, or better yet had never existed at all.

It is true that all of the transfer proposals currently in some stage of development or implementation dictate a minimum size for IPv4 transfers. But the prevailing convention for handling this is to define the smallest permissible IPv4 transfer as equal to whatever minimum allocation size is dictated by concurrent policy. Since transfers themselves have been approved in large measure because the community belatedly determined that community members would be unlikely to adhere to any policy that contravened community member demands, it is reasonable to assume that the policy-defined minimum will also be reduced, progressively, as soon as new demands become apparent.

Of course, it is conceivable that policy will not respond to future demands -- or even that it will respond in the negative for the first time in recent history. For example, if the service providers that voted themselves the freedom to dispose of IPv4 as they see fit subsequently conclude that the resulting fragmentation/inflation of the routing table is happening too quickly, they might seek to impose and enforce the kind of strict policy-based limits that used to exist in the early days of the RIR system.

However, back then, the "safety valve" that made hierarchical routing and route filtering by prefix-length sustainable was the RIR system itself. Assuming they were large enough, aspiring new entrants that wanted to pursue their own independent routing policy always had recourse to the direct RIR allocation process; IPv4 in globally routable quantities was always available to would-be

routing service providers with demonstrable need.

That canonical RIR-era arrangement -- hierarchical routing and prefix-length filtering for the big operators plus a neutral, open allocation mechanism for new entrants -- created an excellent, conflict-minimizing industrial environment for growing the Internet. Aggregation and prefix-length filtering gave large service providers the sustained commercial benefits of economies of scale in routing service provision -- i.e., +++ customers at the price of + routing service provision costs. At the same time, the eligibility-based, RIR-level allocation process provided big startups and successful, growing recent entrants with a transparent mechanism and goal line for becoming "provider independent."

Everybody won, or at least could see clear benefits from the arrangement. In fact, one could probably make the case that the merits of this arrangement contributed to its replication in other regions, ultimately resulting in the current, regionally organized but globally uniform RIR-based system for allocating IP number resources (note: academics like Hendrik Spruyt and Nobel laureate Douglass North have provided similar explanations for how the global political structure ultimately came to be absolutely and homogenously defined by the institution of sovereign nation-states).

Now, however, that bargain has been abandoned -- and it seems highly likely that the party that lost out (i.e., the new entrants) will ultimately react the same way that they do in other industries that are shaped by a concentration of market power/control over critical bottleneck inputs.

So, in the end, we may get to run the market advocates' experiment and see exactly what the routing table looks like when there are no durable limits on demands for finite routing system capacity.

Time will tell, perhaps real soon now...

RFC 1744 – Predictions of the Market to Come

February 1 ***COOK Report***: what happens if the up stream carrier that the community network is deigned to bypass can't get ipv4 addresses that are routable by the upstream carrier? I have Read RFC 1744

QUITE an eye opener!! http://tools.ietf.org/html/rfc1744

Goldstein: Yes, Geoff's paper of 14 years ago was quite prescient. He was making sense in economic terms.

Vest: First, for the sake of others who don't get around to reading RFC1744 themselves, here's a sample of that sense. . . .

"*It is also anticipated that in an unregulated environment the trade in address blocks would very quickly concentrate to a position of address trading between major Internet providers, where a small number of entities would control the majority of the traded volume (market efficiency considerations would imply that traders with large inventories would be more efficient within this trading domain). It is also reasonable to expect that the Internet service providers would dominate this trading area, as they have the greatest level of vested interest in this market resource. This would allow the Internet service provider to operate with a considerably greater degree of confidence in service lifetime expectation, as the service provider would be in the position of price setting of the basic address resource and be able to generate an address pool as a hedge against local address depletion for the provider's client base. There is of course the consequent risk of the natural tendency of these entities forming a trading cartel, establishing a trading monopoly*

position in this space, setting up a formidable barrier against the entry of new service providers in this area of the market. Such a scenario readily admits the position of monopoly- based service price setting. Compounding this is the risk that the providers set up their own "title office", so that in effect the major trading block actually controls the only means of establishing legitimacy of "ownership", which in terms of risk of anti-competitive trading practices is a very seriously damaged outcome."

And some more:

"*The most negative aspect of this are is that these processes will erode levels of confidence in the self regulatory capability of the Internet community, such that significant doubts will be expressed by the larger community the Internet process is one which is appropriate for effective formulation of common administrative policy of one of the core common assets of the Internet*."

"*These outcomes can all be interpreted as policy failure outcomes*."

And finally:

"*It is also appropriate to conclude that continuation of current address space management policies run a very strong risk of restrictive and monopoly-based trading in address space, with consequence of the same trading practices being expressed within the deployed Internet itself."*

There's not a lot about "natural pricing" or rational, sustainable, or otherwise happy outcomes of any kind in RFC1744.

Editor: [And later - *Vest*] -- I'll quote the preamble [in seven parts] to my own (unsuccessful) suggestion for how to deal with the addressing crisis:

Part 1. IPv4 exhaustion is inevitable.

2. Black markets and/or competitive abuses of IPv4 are inevitable, forever, as long as IPv4 is an absolute bottleneck to entering the Internet industry.

Goldstein: So far so good. Black markets happen if gray or white ones can't, of course.

Vest: Black markets also happen when gray or white ones *do, also* happen, of course -- especially when the difference between black and white is determined post-facto, by the willingness of buyer and seller to publicly record their deeds.

Earlier **Vest**: 3. Black markets and/or competitive abuses will inevitably lead to institutional failure (i.e., of current mechanisms of industry self- governance) if not the absolute failure of Internet systems and technology itself.

Goldstein: That's a stretch. Black markets are wonderful lubricants for systems that lack official flexibility.

Vest: The problem is that not all forms of "flexibility" that may be demanded in individual transactions are actually sustainable by all markets. **The demand for confidentiality or non-transparency, and for freedom from "bureaucracy" is common if not universal, and quite understandable in most cases. Whenever transacting parties can get away with non-disclosure, they usually do, for better or worse (e.g., in the banking sector, better yields right to individual privacy, worse yields global financial collapse). In this particular case, the complete absence of any countervailing enforcement mechanisms is sure to encourage more people to press that demand. But as a result, the registration database that is the only mechanism that assures the uniqueness of IP number resources (and**

hence their basic utility) will cease to be sustainable.

Maybe the private sector would evolve alternative mechanisms to mitigate this registration/uniqueness problem, but the result would be equivalent to the creation of mutually independent alt roots in DNS. They can be used to exercise market power too (c.f., China).

Goldstein: They're part of the real world. Competitive abuses are a different issue, but if market power is that concentrated, IP addresses are the least of our worries.

Vest: BINGO! **IP addresses are the least of our worries *today*, because even though market power exists, it is exercised by means of other critical, non-substitutable bottleneck inputs (e.g., last-mile facilities).**

IP addresses are [also] the least of our worries *today*, because the exercise of market power via IP addressing was maintained at (generally acceptable) low levels, by the operation of the RIR system -- that is to say, by the continued existence of a supply of the critical, non-substitutable bottleneck input that could be obtained "as necessary" at fixed prices from a non-competitor. I really don't think we want to say that the existence of market power in one product market is undeniable, so therefore all efforts to prevent the emergence of market power should be abandoned everywhere, because monopolies "don't matter anymore" - i.e., this is already the worst of all possible worlds….

Earlier **Vest** 4. There is zero possibility that IPv4 "resource transfers" or official markets will prevent black markets and competitive abuses of IPv4, or the failures(s) that will result.

Goldstein: Whoa. Too complex a sentence. White-market official markets will reduce demand for black markets. Either way, markets will exists. Market failure happens when power is excessively concentrated, which implies difficulty in market entry.

Vest: I guess I subscribe to a broader understanding of what "market failure" means. To me, it means that markets fail, period. For example, I interpret the ongoing global financial collapse as a market failure, even though market entry restrictions didn't have anything to do with it at all.

Earlier Vest: The only way to avert that outcome is to remedy the root cause, which is the non- substitutability of IPv4.

Goldstein: I wouldn't say "only way", but it's one way. Reducing the scope for which IPv4 is non-substitutable would also reduce demand and thus price.

Vest: We are agreeing here I think, you are just focusing on reducing a specific dimension of non-substitutability, whereas I was describing a general, qualitative reduction in bottleneck status.

Goldstein: NAT is one such approach. Corporate nets often substitute private for public address space, even if just for security reasons.

Vest: NAT and non-unique addressing is one candidate approach for everyone *except* aspiring routing services providers. If NAT and non-unique addressing was an equally acceptable option for aspiring routing service providers, then this conversation would never have taken place.

Earlier **Vest** 5. IPv6 provides one possible means for solving that root problem.

Goldstein: We disagree. IPv6 is so broken that it doesn't help; it was a huge error. Thus I suggest we seek better alternatives and stop wasting time on IPv6.

Vest: Fair enough. We disagree on this point.

Earlier Vest 6. However, in the presence of radical uncertainty between competitors, there is no high probability path leading to the timely widespread adoption of IPv6 that does not involve some level of inter- provider coordination.

Goldstein: Well, IPv6 ain't happening, but inter-provider coordination is what the Internet (as a business model, not a protocol suite) is about.

Vest: Okay.

Goldstein: How much coordination is needed depends on the technology chosen to implement the model.

Vest: That seems like a logical assertion, but I think the implication that you want to make is actually an empirical question that remains to be answered**. I know that Patterns of Network Architecture (PNA) by John Day aspires to change (or rather eliminate) the scope of required global, a priori coordination (or standards sharing) in some ways. In some respects the LISP protocol development effort is attempting to provide a similar, if less revolutionary mechanism to the same end. But the devil is always in the details, and the symmetry between addressing/routing and money/financial flows has left me with the intuition that the closer you look, the more likely you are to see variants of the old familiar devils.**

Maybe PNA will be different, or maybe not; we'll have to wait and see...

Editor's Postscript: A somewhat revised and definitely more polished version of my summary essay on the risks of IPv4 from the March 2009 issue - last month- is found on Circle ID at http://www.circleid.com/posts/20090202_ipv4_numbers_transferable_property/

KPN Joins Amsterdam in Extending its Open FttH Network

Paul Budde: Today (February 4, 2009) the City of Amsterdam announced its move to the next stage of their FttH project - with another roll out covering 100,000 connections - is a clear indication that the concept of open access FttH networks is a valid one. This will have large scale implications for countries around the world who are looking at using open network based telecoms infrastructure projects to stimulate their economies.

These developments are important as it validates visions such as presented in the reports we have prepared for governments in the USA, Europe, Australia and New Zealand and start answering the question s asked by policy makers whether this is achievable.

It is a clear indication that open networks are the way of the future. They will deliver unprecedented economic and social benefits. This will force government and industry to move into a trans-sectoral mode of thinking – this time with an incumbent eager to promote that message as it now clearly recognises the benefits that can be gained from such an approach.

For more info see: http://www.buddeblog.com.au/open-networks-delivering-the-goods-for-all/

Editor: This link is exceptionally good. Here is the introduction:

Paul Budde writes: *The City of Amsterdam announcement to now move to the next stage of their FttH project - with another roll out covering 100,000 connections - is a clear indication that the concept of open access FttH networks is a valid one. This will have large-scale implications for countries around the world that are looking at using open network based telecoms infrastructure projects to stimulate their economies.*

BuddeComm has been involved in industry policy discussions with experts on three continents about the future of telecoms.

We have developed a scenario. The vision naturally has a range of strategies attached, but in simple terms it works as follows:

** Telecoms infrastructure is of national economic and social importance;*

** For our societies to profit from the digital economy infrastructure must be based on the principle of open networks;*

** This allows us to multiply the benefits this infrastructure has to offer to other sectors such as healthcare, education, energy, environment, media and communications;*

** Once open networks and the access tariffs are established the national telecoms operator will be in the best position to run this network, thus avoiding the necessity for wasteful duplication.*

** Infrastructure and digital applications will need to be developed parallel with each other and this requires trans-sectoral thinking from the government and the industry - not the current silo thinking.*

While it is great to put such visions in front of the policy-makers in the end the question is whether this is achievable - can such a vision be implemented?

Again the Netherlands is paving the way. Amsterdam was one of the first to identify the multiplier effect and the need for trans-sectoral thinking. Unfortunately the incumbent KPN first preferred to roll out fibre nationally together with the cable companies and so

to hold on to its vertical monopoly. Only in the last few years KPN started to warm up to becoming involved on the basis of an open network.

Amsterdam fought legal battles against the incumbent cable company UPC, a property of Liberty Global's tycoon John Malone, in both the Netherlands and European courts. The City won this battle and a consortium, which did not include the incumbent, started to roll out the city's FttH network.

This development became a real boost for other FttH networks in the Netherlands, and they favoured the independent fibre builder Reggefiber.

[SNIP] **Editor**: Much more valuable reading on open networks at the above link with additional valuable pointers to open networkselse where in the world.

Felten: Paul beat me to the punch, but you can see my positions on this here:

http://www.fiberevolution.com/2009/02/amsterdam-phase-deux.html

The Official Amsterdam Announcement

Dirk van der Woude: At last I can reveal what we have been cooking in Amsterdam.

City of *Amsterdam* - Press Release
Amsterdam, 4th February 2009 11:15 CET

*Further roll out of open access fiber network in ***Amsterdam*** with
Reggefiber / ***KPN***

Mid 2009 Amsterdam will see the next stage of the roll out, covering another 100,000 open access fiber connections. This follows from an agreement between the initial owners of Fibernet Amsterdam and the joint venture of Reggefiber and Dutch incumbent telco KPN. The agreement ensures the construction of a high quality and future proof communication infrastructure, a basic condition for the city's economic and social prosperity.

The contract ensures the open access character of the network of now 43,000 homes connected and passed. Other service providers or operators are as welcome as before. The agreement will now be submitted to the Dutch competition authority NMa.

Every year the need for a faster network increases. In just 15 years the percentage of internet users in the Netherlands grew from zero to today's 90%-plus, the large majority of them broadband users. Alderman Van Poelgeest: "Fast networks are important for the future of the city of Amsterdam. We want our citizens to be offered the best in telecare, e-Health, distance learning and teleworking. As the construction of this fiber network will take quite a few years the right time to start is now."

In the Amsterdam network every address is connected directly with its own fiber to the neighborhood switch house. This 'point to point' architecture ensures capacity even when many neighbors are teleworking or watching videos at the same time. The environment benefits as well: fiber only networks consume 10 to 15 times less energy than networks that combine glass fiber and copper.

OPTA / NMa

The open access nature of the network is fully in line with last December's decisions of the Dutch telecom regulator OPTA and the competition authority NMa on glass fiber networks as well as the cooperation between Reggefiber and KPN. Alderman Van Poelgeest: "Amsterdam principles meet the conditions of OPTA and NMa and have KPN's full support. Therefore, I am very pleased with this deal. The Amsterdam municipality keeps their promise: further roll out of

the fiber network, together with strong and determined market parties. Municipal involvement now can decrease, probably eventually to zero. Last but not least, in the view of the current circumstances this fiber deployment at once generates employment and it boosts the Amsterdam economy." The participating parties focus on a fiber roll out in an economically viable manner. This means that after or during the roll out to the next 100,000 lines a decision will be taken on the then remaining 250,000 Amsterdam addresses.

Decreased shares for housing corporations and municipality, *KPN* service provider

In the initial design the four housing corporations (Ymere, Stadgenoot, Rochdale and De Key) jointly had 33% of the shares just like the municipality and Reggefiber. In the new situation Reggefiber through additional deposits acquires 70% ownership. The municipality and joint housing corporations keep 30% ownership. In the agreement it is stipulated that key decisions can only be taken with an 80% majority vote. An independent 'Priority Foundation' will hold a golden share with veto rights for key issues such as the open access nature of the network.

Starting Autumn 2009 KPN will offer services on the new as well as existing parts of the network, in open competition with other service providers, in conformity to OPTA / NMa rules and regulated rates.

Historic Step

This morning's signed agreement is the first of its kind: until now there has never been an example of this kind of cooperation to expand an open access fiber network. Amsterdam's Mayor Cohen: "I expect the expansion of the open fiber network to have far-reaching positive implications for Amsterdam's development. Today, like energy and water supply broadband is an essential necessity that should be accessible to all.

Vincent Dekker: I have an interview with KPN CEO Ad Scheepbouwer on open networks in Thursday's *Trouw*. It shows pretty well how much he is in favor of open networks.

[**Editor**- Dutch version at http://www.trouw.nl/nieuws/economie/article2023373.ece/KPN__Glasvezel_delen_is_slimmer_.html]

Some quotes:

'In hindsight KPN made a mistake back n 1996. We did not react too enthusiastically to the obligation of allowing competitors on our old wireline network. That turned out to be not very wise. **If you allow all your competitors on your network, all services will come on your net, and that results in the lowest cost possible per service. Which in turn attracts more customers for those services, so your network grows much faster. An open network is not charity from our side, in the long run it simply works best for everybody.'**

'As far as KPN is concerned, services are the future of telecom. We're happy to invest in the network, but we don't want to be only bit sellers. The market for services is much bigger.'

KPN-people acknowledged figures Tim Poulus picked up earlier in Rotterdam: http://telcommunicator.blogspot.com/

- after losing subs for years to cable, in the first few towns where KPN now actively sells ftth it is regaining market share.
- it's ARPU for ftth is 58 euro excluding VAT, where ARPU for old ADSL plus VoIP is only 30 Euros. part of the extra 28 Euro is for TV, which KPN could not really offer on ADSL.
- penetration is growing rapidly now that KPN is actively

selling ftth. after ten weeks or so it now stands at 20 to 30 percent, according to KPN.

Penetration is the number of homes that really use ftth as a percentage of homes passed/connected in a house to house roll-out.

Remember that Reggefiber rolled out ftth in two ways
- in some towns roll out only started if 40 percent had signed up. here penetration is now growing to 70 percent and more, according to Reggefiber.
- in other towns it rolled out ftth without prior sign-ups. this turned out to be a tougher sell.

But now that KPN has joined Reggefiber it is far easier. KPN can promise it's pots-subscribers a switch to ftth without any hassle. 'Nothing changes, you only get much faster internet, free phone calls to the entire country and much better picture quality for your TV. And it will cost you about the same or even less then top speed ADSL-internet, phone and TV-from-cable combined now cost you.' KPN has taken it's time to get everything worked out. I think they're ready now and the next few months must prove that ftth is an easy sell...

Scheepbouwer also acknowledged some arithmetic of mine.
For those of you who think ftth is expensive for KPN, think again:

Rolling out ftth costs on average 1000 euro per home passed. Ftth for the whole of the Netherlands will need some 600.000 homes connected per year, which foots a bill of 600 million a year. But the joint venture of KPN and Reggefiber will get 60 percent, 360 million a year, from bank loans and 40 percent, 240 million, from the shareholders. KPN owns 41 percent of the joint venture so has to put up 100 million euro a year.

For a company that has a free cash flow of 2.5 billion a year that is not too big a problem ... "

IT, Internet and the Power Grid

Ken Miller: *Note: The executive summary of the study is currently available at http://www.jcspstudy.org. The full study will be posted within the next several days.*

Study: Billions Needed to Deliver Wind Power to Eastern Interconnection

Joint Coordinated System Plan Estimates $80 Billion in Additional Transmission Needed to Deliver 20% Wind Energy

(CARMEL, Ind.) The Joint Coordinated System Plan (JCSP'08), the first step of a transmission and generation system expansion analysis of the majority of the Eastern Interconnection, estimates the electricity sector will need over $80 billion in new transmission infrastructure to obtain 20% of the region's electricity from wind generation. Snip

St Arnaud: And this is why we have a program to move computers and data centers to wind locations, rather than depending on the grid to bring renewable power to the city. We have several such projects under development Not only do you get cheaper power at guaranteed price, the electrical line transmission losses that would incur in bringing the power to the city pay for the fiber to the wind energy site

http://green-broadband.blogspot.com/

Ed Pimentel:
http://www.google.org/powermeter/howitworks.html
http://www.google.org/powermeter/smarterpower.html

I am seeing a number of new business opportrunities here. Think of a Zigbee, Bluetooth2.x, Wifi device that reads the SmartPower meter info and sends to Florida Flickr & Flash power company. Or uses this to provide Co2 credit to biz or consumers... Think kitchen appliances..... services...

Van der Berg: Google is not that orginal ;-) These Brits have a very beautiful solution already
http://www.diykyoto.com/ It changes color to indicate how your doing with your usage and has digital numbers to tell you more. They use software called Holmes for the analysis.

Coluccio: Since no one responded to an earlier message of mine here several days ago, concerning the use of swarm logic applied across a wireless mesh network supported by ZigBee, to effectively achieve many of the same ends now being discussed, I'm re-posting it here for the benefit of those who may overlooked or not received it:

Managing Energy with SWARM Logic: Self-organizing equipment could cut energy bills. By Tyler Hamilton | February 4, 2009 *MIT Technology Review*
http://www.technologyreview.com/energy/22066/?nlid=1762

For the power industry mavens on the list and others who are astute in this space, a couple of questions:

Assuming the use of the technologies being referenced here today can effectively measure, and thereby manage, loads, how do these approaches impact the business models keyed to the arbitrage of demand-response?

Secondly, What will be the effect on the demand-response model resulting from the assignment and future trading of carbon credits?

All thoughts on these aspects of the approaches being discussed here today are welcome.

Miller: I saw the post but have been buried. This is an interesting concept, but I am uncertain as to the degree of effectiveness of the claims. "The devices learn the power cycles of each appliance and reconfigure them to maximize collective efficiency."

This will keep the heavy equipment from turning on at the same time (I would bet a very similar concept to CSMA/CD), but at the end of the day, will it reduce the total power consumption? In my experience with the buildings and systems that I have worked with, it will not. I am not saying it would not be useful, but time shifting consumption a few minutes earlier or later has little affect on total power consumption. You may be able to reduce your peak demand, which depending on your agreement with your local utility will set your bill higher, but it wont affect consumption. Generally in large facilities this statistically averages out anyway.

Now, if you time shift to alternate times of the day, then you could take advantage of rate differences based on time of day. The vast majority of people cannot or do not have access to a variable rate agreements, unless you are of significant size and sophistication. The organizations that are big enough generally do manage this because there is a clear ROI.

For must facilities, the user comfort or business need is generally a just in time fulfillment process. Unless you have the thermal or energy storage "systems" (read capital and designed in up front) to keep people and things cool during the day without consuming power, this is not something a swarm can help with.

Also, the product claims that "Before making a decision, he explains, a node will consider the circumstances of other nodes in its network." In most cases buildings already have controls that all come back to a central system. The issue is lack of software or "global" system thinking not the lack of controls. I love the idea of the "swarm", but how about putting your algorithms into software that is already deployed in building systems to affect these kinds of behaviors. Most building control systems I am familiar with could do this tomorrow with the installed hardware. The challenge is programming the logic.

On a side note, is it scientifically possible for swarms to globally optimize? Certainly they can do better than most systems because most systems don't try any optimization of this kind. From an AI experience there are many many advantages to self-organizing systems, but I don't remember that they regularly reached globally optimum solutions.

All that cynicism aside, if you were doing a Greenfield build AND you could deploy all wireless swarm controls **instead** of network of controls wired back to centralized controllers, AND you have management software that could plug into the swam to scope and shape behavior, AND they do everything the old controls do, AND they can reduce the power bill, AND they don't cost 100X the standard controls in the install and life cycle, this starts to get real interesting.

I really am starting to sound like a Luddite….. While I am really for these kinds of things, I am just seeing a great deal of marketing promises that I would like to see better data on. The key for me would be very inexpensive (?), secure (Zibee is working on that), automatically mesh themselves (REGEN appears to be there), and management software to track, monitor, and manage.

Coluccio: Thanks, Ken.

Individually, these new tools may impact very little on the larger front, but the trajectory that they suggest is unmistakable.

Miller: Absolutely. These building controls have

evolved very little over the last 10-20 years. They are difficult to program and lack any intelligence what so ever.

I wonder if this evolution will be similar to APC's approach with UPSs, PDUs, Rack PDUs, and CRACs. APC began embedding Ethernet/IP controllers into their equipment. It has the capability to start with SNMP, but has direct HTTP configurability, and some proprietary controls mechanisms (?soap/xml?). Most installations use only a fraction of the capability, however, as more and more equipment comes enabled, the value of this "management network" goes up (N^2 Metcalfe?).

Having slept on the swarm idea for these controllers, I get more interested. Not so much for the marketing reasons in that press release, but what could be done from a practical configuration and programming level. I am still skeptical of the energy savings of such an approach, but a mesh and swarm could have a big impact on setup, maintenance, and capability of a building control system.

I think that I am going to get a few to play with.

St. Arnaud: I agree fully with Ken Miller's points.

Most people think today's Smart Grids will reduce consumption- where in fact they largely only displace consumption from peak demand. In the several studies done s far on smart grids and smart meters, the savings have been very small for consumers.

Some early studies indicate that providing feedback to consumers on energy consumption will reduce demand. But these studies are very limited scope and the sample population is usually made of dedicated and concerned homeowners. We will need longer-term studies with larger populations to see if feedback meters will have a meaningful impact.

The biggest beneficiary of smart grids is not the consumer, but the utility who does not have to build more power plants for peak load. Indirectly this marginally does reduce CO2 emissions because base load is usually nuclear while peak load is usually gas powered plants

Coluccio: Bill, I see the merit behind your comments and I agree that Ken makes a convincing argument as well. However, one can't treat the technologies related to electric power grids (the emphasis here is on monitoring and control) that have brought us to this point in the discussion as though they were universally applicable across all measures of scale. Perhaps telecom networks lend themselves very nicely (or more favorably than do power utility networks) to fractal analysis and for extrapolatory purposes, but when it comes to the economics underlying the surveillance, monitoring and control of electric power distribution systems, the same cannot be said in as sweeping a way.

The swarm algorithm approach mentioned earlier, for example, despite its stated purposes in the marketing materials I posted, may actually be very beneficial in assisting facilities managers of large buildings and campuses in the administration of building automation systems, yet they may well have little if any influence at all in effecting cost savings on the utility bill. I dare note, however, that if a demand-response arbitrage operator is capable of satisfactorily reducing peak load utilization, hence the costs to the customer as well, by placing a telephone call to a building or campus owner and advising them to shut down their discretionary demand, then I submit that a swarm device capable of sensing everything that the demand-response operator can sense, could very well do the same.

Budde: I disagree guys.

In my work with smart grids I

do get overwhelming support from customers, simply for the fact that a smart grid with a smart meter for the first time in their lives gives them an opportunity to see what is happening with energy in their house. Simply the fact of knowing this and that fact that many people are motivated to better manage their energy will see changes happening in their behaviour. You give them a tool that they can use to do their bit for the environment; there are millions of people who would use such tools and will use people power to start making changes. Keep these people dumb and they are powerless to do something

It is similar to broadband, what would be the benefit of broadband to people say 10 years ago? With all the applications that since have been developed broadband has now become an essential service to many people.

On the other hand most utilities make money from selling more and more energy, smart grids will reduce energy consumption by 25%, through efficiency gains, without having an effect on life style and they will feel that in their bottom line. So what is the benefit here? (Don't start throwing stones as I deliberately put this in black and white)

The generators are indeed the companies that profit as they don't have to build these new stations.

I realise that the US is running behind as you don't yet have a universally fully structural separated energy industry, this leads to several conflicting elements within this industry sector which in turn can lead to a polluted vision on smart grids.

Sterling: I have to agree with Paul and Frank on this matter.

The bottom line is that (let me preface by saying I'm talking about USA policy) regulators at the state and federal level have basically given electric utilities a free ride for years in terms of getting a regulated return on supply side investments.
Smart grid in conceptual terms has been around for at least 10+ years. BPA calls it "Energy Web" , EPRI calls it "Intelligrid", and many joined a group called "Gridwise".

The basic premise is that a "smart grid" can be created out of our existing grid if we replace outdated analog infrastructure with digital infrastructure. The new grid needs to have energy storage, distributed generation, DSM (demand side management), and the highest level of energy efficiency that is possible.

The best kept secret in the electricity business is that the levelized cost of efficiency improvements is way lower than any other possible resource available. Plus it is a firm load. There are ESCOs (energy service companies) that make a handsome living going after efficiency projects, but negawatts have always been the Rodney Dangerfield of the utility industry.

Supply side rules, demand side drools.

Smart grid could provide the opportunity for an entire community or city to be considered a master metered campus which ought to be treated as a negawatt power plant. Using an ESCO model, that community should be compensated for verifiable energy savings just like a wind farm is compensated.

Under this model, what's in it for the electric utility? Well that is the big question. Electric utilities are between a rock and a hard place. I am sure that most would prefer to continue the status quo but between cap & trade, climate change, and renewable energy portfolios the utilities are facing the certainty of change.

The debate over smart grid covers a spectrum of solu-

tions that on one end may be the equivalent of open access or structural separation vs a walled garden or smart grid lite on the other.

None of these smart grid conversations pertain to the transmission capacity issues described by T Boone Pickens or other Big Wind proponents with one exception. Some of us believe that each community needs a resilient smart microgrid with all the efficiency, conservation, chp (combined heat and power) and local renewables prior to seeking outside supply. Utilities and other big power players want the top down planning described in Frank's previous email.

This fight is at least 30 years old. I participated in the old fight during the WPPSS debacle (the largest bond default in the history of municipal finance) in the late 70s.

It may only shift demand in time, but that's GREAT. Moving demand from periods of high demand or tight supply (logically equivalent) to periods of low demand or slack supply is a lot like storing electricity, with the added detail that it's the only kind of storage that will ever be 100% efficient. It's not often you get to buck entropy.

There are some interesting experiments going on in Europe with this; for example, the Dutch have a huge amount of cold-storage warehousing on the port of Rotterdam and elsewhere. This is nice because the stuff in the warehouse doesn't mind being frozen colder than it is already. They also have a lot (15% or so) of wind power, which fluctuates.

So they have a trial going on where the refrigeration plant is hooked up to a feed of updates from the grid reliability control.

When the wind blows, and it isn't load-following, i.e. there is a surplus of power, the refrigeration cranks up and cools the warehouse down a few degrees, taking advantage of the cheap electricity and also relieving the grid of it. Then, when the wind drops or more load comes on line, the refrigeration shuts down, taking its demand off the grid and letting the warehouse gradually warm up towards the normal target temperature. Power is effectively stored in the sides of beef or cheeses or whatever.

Elegant, but I don't see the relevance of "broadband". There's nothing broadband about signaling and status messages - it's all very much like SS7, SMS, XMPP, SIP or whatever push-messaging protocol you like.

Cole: The "relevance of broadband" is that the tiny signals involved for the smartgrid are, more often than not, sent via optical fiber because (a) often cheaper than copper; and (b) can be placed close to high-power lines without electrical interference. Thus, with a SEPARATE lambda, that same fiber can provide huge data capacity for the rest of us, without compromising the security of the electrical control, etc.

A number of municipal electrical utilities have already discovered this, and are finding they can "share" the fiber they were installing anyway. Plus, they have a roughly similar customer facing team (billing, etc.) that adopts relatively easily to dealing with triple-play (and more) customers.

Miller: The following monolog is my opinion and not representative of the organization or industry that I work it. It is also a quick dump of some thoughts, with only minor edits, so I apologize in advance for incomplete sentences (and thoughts).

When Smart Grid comes up, I think a great deal of confusion is coming out of differences in terminology and mixing concepts. I think there are very large differences in what SmartGrid means to everyone and therefore very different view-

points, leaving a lot of room for each perspective to have some solid truths, but too much room to disagree.

I think confusion comes from over generalizing the following three major conceptual groups:
1) The application of "Smart" or "Intelligence" to local "distribution" vs. "transmission" vs. "load".
2) Concepts of efficiency, demand reduction, and peak shifting.
3) Complex issues with Generation sources such as environmental friendliness, use in base load vs. ramp vs. peak loads, speeds to ramp up and ramp down, reliability, and cost . These are all very different issues that are interdependent, but most often just get lumped together in conversation which creates a great deal of confusion.

Smart intelligence
1) Generation, Transmission, and Load Regarding "intelligence", the transmission system and dispatch of generation is already what most people would call smart. At the Midwest ISO for our footprint (first or second largest power footprint in the world), we monitor all power flow and do full-state estimation for a 600k point network model and dispatch generation for all 2500+ generation points in the footprint on a 4 sec/4min basis. We control the frequency, demand response, of the generators, etc. All in real-time (if 4 second basis is considered real-time). Generation to fulfill load is pulled from power markets which are bid into in advance and the most economic solution to match load is based on market prices. Also, the reliability of this large scale system is very high, other than some notable blackouts.

Unless I am wrong, most SmartGrid discussion is not talking about Generation and Transmission.

Most comments about the lack of intelligence and reliability are focused on local distribution. Without a doubt, there is a lack of information and ability to reroute around problems. This can definitely benefit from better controls, but will also need significant topology modification to improve reliability. This is very capital intensive and has little to do with Smart Meters and data enabling distribution substations and transformers.

So when the there are claims for "more efficient", what are people talking about? I think most people mean that all the "lost power in the sky" will become available to be used. Losses in the electric grid are largely at the Generation and the Load End not in the grid. Transmission (I will need to look up the stat to verify) is less than 5% of energy loss from source to load. Local distribution losses are different and higher (there is some disagreement on the amount of loss in local distribution), but there is no doubt that the reliability is significantly lower in distribution the routes are significantly less redundant (lack of multiple sources to feed large segments of distribution). Topology is like sonet, but most of the local distribution network links are tributaries.

So if Smart Grid local distribution does not deliver massive amounts of additional power, what and who benefits from smarter distribution? Consumer gets increased reliability and the ability for local power injection. I think this is where many smart grid advocates see advantages. There is a belief that local generation will be far more efficient with local injection. I am unconvinced yet on this given that most losses are NOT in transmission and distribution (I am certainly not against this, and think it could work). Additionally there is a belief the local generation with SmartGrid could automatically balance. Having seen the math and computational power for this to work, I am interested but still on the fence for this particular aspect. (It could work well).

SmartGrid in this DISTRIBUTION concept is referring to

the adaptability and performance of the infrastructure and having little to do with reducing the load.

Demand Reduction

2) Efficiency, Demand Response/Reduction, Peak

Another common statement generalized in SmartGrid discussion is "Power Savings" or "Efficiency". Most people are lumping too many different concepts into the term "Efficiency". Although, semantically, you might be able to make the case it is correctly used, the term is being used by people to mean so many things, that arguments often ensue.

Technically, making loads more efficient is about one time improvement, not Smart or adaptive. An efficient motor is more efficient when installed or upgraded, then it saves over its entire life. This area is by FAR the largest area for improvement in reduction of power consumption, but it does not need a single Smart Grid component (the device itself might be smart to operate efficiently, but having nothing to do with the grid). Personally this is a belief and passion. Most of power is wasted by the way it is consumed at the load. There are many many ways to improve the grid, but addressing the vast installed base of motors and compressors that struggle to get 40-60% efficient is big (when they can be 90%+). Personally I believe that the entire smart grid in the US over the next decade will not reduce the power consumption the way that efficient equipment at the load will reduce it. Over the longer haul the SmartGrid will be needed for other important reasons such as more reliability and information source for load decisions. I am concerned that the smart grid emphasis will take our eyes off of making the appliances and motors efficient.

SmartGrid (as an information source) comes into play if a motor responds to increasing power price or request to reduce demand and stops its work. This is Demand Response. Consumers and businesses can make decisions to not consume power at certain levels. I believe this is what most people think of when SmartGrid and Demand reduction is talked about. The challenge with this concept is that in practice most activities cannot permanently stop the work. It can only be delayed. Some people argue that demand response will reduce environmental impact or reduce power consumption. I think this is misleading, because, most often the load is only delayed.

If it is only delayed, this is really not demand reduction but Peak Shifting. It is still a valuable tool, but I question calling this efficiency. You might call it "efficient use of the available capacity of the grid", because this type of response can reduce peak load, but it is generally lumped into the concept of demand response. This type of "efficiency" is not what most people understand that SmartGrid will bring about. At the end of the day, with the same generation mix, moderates CO2 production through more predictable consumption, but does not really reduce demand or make the Grid more "efficient" in the delivery of power. You might call this efficient operation of local distribution, but it does not directly save power.

This is where the "Smart Meter" is normally discussed. It's a good information source for consumer and bill provider, but does little for distribution management (unless we give the local power utility the right to start shutting off things in our houses and businesses without discussion). Distribution management does not need individual meter loads to understand the load on any particular power segment. This can be done at the transformer and distribution substations.

So what does Peak load shifting do (from my wholesale global viewpoint)? It evens

out the power load. This predictability reduces the amount of reserves that must be maintained (which decreases the cost of delivery). It also allows the maximum use of lower cost base load options, such as Nuclear and hydro. A Nuclear plant can take weeks to ramp up. If the load is steady this can be matched effectively every day, maximizing the consumption of the cheapest generation sources. It does reduce additional generation construction of Generation, but only for a while while we take the in-efficiencies out of the system. Then when growth returns, additional generation will be required.

This is a much bigger topic and I am stopping short of raising a bunch of issues related to this, but I wanted to point out that Smart Grid as an INFORMATION SOURCE will _primarily_ allow this peak shifting, but in my personal opinion, little net power demand reduction. Unless it stimulates consumers to buy "efficient" devices (human behavior not technological).

Smart Generation
3) Generation

As we start talking about Generation I must give the caveat that I am around industry people all the time, so my view may be colored.

To my understanding of Smart Grid, it affects generation only in the sense that you can inject local generation. All the same economic and environmental issues associated with different types of generation don't appear to be affected. · Most of the benefits of reduced demand (through real efficiency), I argue do not require Smart Grid. · Peak shifting is largely a function of Smart Grid as a Peak Shift/demand response.

Personally, I am all for local and distributed generation. I have done this personally and been involved in these types of projects before. BUT, but there is a long way to go for this to work with SmartGrid. (I would like to keep the technical issues separated from the legislative discussion, which is an entirely different barrier to this concept). · Safety is a big issue. Smart Grid can help this with the rapid and proper isolation of faults. · Reliability. Everyone wants wind and solar, but when it stops, they want power from somewhere else to be there! This requires significant advancement in energy storage approaches. I don't get yet, how smart grid helps this. Perhaps this tells people when to buy from the grid and store vs. use their own generation. · Balancing load across a grid is very difficult. How does smart grid do this over a large scale? Its easy to talk about small pockets keeping themselves covered, but when faults occur and generation sources go offline, there can be HUGE power draws from non-local sources. Today, these small generation sources just isolate themselves to protect their community. This immediately causes other issues… and on and on. I am open to understanding this, but I have yet to see a technical/ engineering description of how this would work. Again this is a big topic, and I stop short intentionally.

St. Arnaud: Excellent summary. The "average" transmission lines loss in the high voltage transmission systems is 7.8%. DC systems do a little better. Many factors such as corona discharge, load factor etc can affect transmission lines losses

But I agree with you the challenge is in all the inefficient motors and AC/DC converters

It all Depends

Coluccio: Hi Jeff.

I think you've aptly demonstrated the high degree of variability between disparate types of situational types (national, regional, local, campus, enterprise, residential), hence the dissimilarity of circumstances and solutions called for by each of those situations as well.

And yet, we often use the term 'grid' to apply to them all as though the meaning of the word is constant. The act of collectively dissecting this issue is but one of many examples why Gordon's reply to Ken regarding the applicability of discussions about the power sector is not only relevant from an interdisciplinary perspective, but also because it illustrates how the power utility sector closely mimics the planning of Internet and telecom plant (including the mediation functions required by each) as well.

Consider the broadband situation in which we in New York City now find ourselves. I am only now, after joining the NYCFF list, beginning to take the time to seriously grock the many sides of the issues that prevail in NY and other large cities that appear upon first-blush to be "over-supplied" markets. Incidentally, and speaking of trans-sectoral exchanges, I'd like to thank Paul Budde for taking the time yesterday to post his best wishes to the NYCFF list and for imparting a generous dose of his wisdom there as well. Thanks Paul!

What I've stated earlier about power grids I believe also applies to the situations that citizens across all sizes and types of locations face when choosing an optimal approach to provisioning broadband as well, specifically referring to the Day 1 criteria that city planners and grassroots movements alike should take into account while formulating a consensus on what needs to be done and how best to proceed.

Unspoken in most discussions concerning the need for public networking initiatives (no matter whether municipally-motivated Institutional Nets (I-Nets) or grassroots neighborhood nets) are the influences on citizens' outlooks of an over-abundance of "triple-play" choices, Instead, virtually all approaches to assessing broadband shortcomings at any level have result from an almost-singular focus on quite the opposite dynamic: Scarcity.

Scarcity is seen as the main driver behind public action. And to be logical about this, scarcity produces a higher degree of exigency than the dilemma surrounding too much choice, obviously, but the examining what happens when a population experiences broadband overload is instructive, nonetheless. Last week I mused to Susan, who is attempting to devise a plan to overcome extreme scarcity that her framing of the problem in some ways was indicative of Stockholm Syndrome. Little did I realize when I exchange notes with Susan last week that I'd later go on to dissect the complacency that characterizes my own area and conclude that the same psychological effect prevails here in NY City as well. Only in the case of NY and other large cities 'enjoying' over-abundance the phenomenon is manifest in tacit acceptance, if not also a rousing level of endorsement of that which was decided exclusively by the incumbent as well.

The similarity here rests in the high degree of reluctance in both camps to define what types of attributes and capabilities their connectedness should possess when the latter are provisioned by traditional providers. Instead, both camps are likely to accept the connection attributes designed and deemed sufficient by the incumbents, who then add on top of those connections 'services' that users may or may not want, or that users may want, but not provisioned by the local incumbent... while citing status-quo indicators usually measured by arcane standards across broader markets when deciding how much bandwidth we 'deserve'.

Returning to the scarcity-abundance dichotomy that Susan and I face, respectively, the differences in our approaches (assuming both causes are deemed legitimate, given that many will suggest the NYC cause is questionable at best), must

necessarily be reflected in two different roadmaps for reaching our respective goals. One might envisage scarcity, due to onerous funding demands on only a few stakeholders, might be best-suited by a top-down approach leveraging a natural monopoly-like form of architecture. In the land of plenty, on the other hand, the unwieldiness of agitating enough citizen support to effect a top down solution is nil to non-existent, given human nature and the formidability of incumbents when they perch as defenders. So, for pockets of users in the land of plenty, perhaps in neighborhoods where higher levels of motivation and enlightenment exist, or within housing complexes and special communities of interest (think small enterprise or neighborhood improvement districts comprised of SMBs) where monetary savings can occur due to economies of scale, neighborhood networks and condo builds that tie into the 'Net may be seen as more appealing.

While I may appear to have strayed quite far from the original discussion, does what I've written immediately above appear all that much different from the situation we were discussing concerning "grids"?

Sterling: Ken, Frank, Bill and others,

I like Ken's structuring of the issue. I'd like to offer three points of my own to kick the ball down the road.

1) Smart Grid is a Vision. As we speak, smart grid doesn't really exist and it's probably an evolutionary process. The question is whether smart grid is a transformative conversation with transectoral implications that goes beyond the boundaries of the IEEE or other private industry group that are looking to optimize the electric utility system. I would argue that smart grid is the connective tissue we need to enable the new energy sector. Simply measuring kilowatt-hours or of peak load misses the point much like measuring megabits per second may miss the point in benchmarking the quality of an Internet experience. We need to create an energy service metric that lets unconventional methods (the negawatt power plant, etc) that meet the needs of the end use compete with supply side solutions.

2) Smart Grid is an operating system. The bottom line for any successful implementation of smart grid is software based. There are firms out there that want to be the Microsoft of the Grid. Google, Cisco, IBM, Microsoft, and a host of start-ups are all trying to figure out what to do. What is the model? I would hope for a open standards approach that would create a development platform for localized energy projects. I'm holding my breath. ;-)

3) Smart Grid is a community network. Smart grid is the ultimate opportunity for transectoral integration at a community level. Unless you include the human factor into the system you are reinforcing the silo mentality that is ossifying innovation in community infrastructure. Take water treatment and reuse as an example. I know of a firm that has perfected an onsite wastewater treatment and recycling system using blackwater composting toilets and graywater recycling to drinking water supplemented by a rainwater cistern. The key to the system is a persistent connection to the Internet (smart grid/machine-to-machine) to monitor water conditions. That solution requires 99.999(9) power and communications, plus a total paradigm shift from both the water company and sewer company. I think people are ready for a change to sustainable and resilient 21st century community infrastructure but siloed utility interests and "well meaning (not)" regulators stand in the way of progress.

Kevin Moss: I think Ken's

summary is excellent. Certainly increased my appreciation of the issues and relative opportunities. Against Ken's framework, it is interesting to note that The SMART 2020 report http://www.smart2020.org/ released by The Climate Group and GeSi towards mid last year identified five primary categories through which global emissions could be reduced by 7.8 GtCo2 by 2020 (representing 15% of global emissions) enabled by ICT services. The five categories and the proportion of reduction potential they represent in the report are

SMART Logistics (transport) - 19% SMART Buildings - 21% SMART Grid - 26% SMART Motors - 12% Dematerialization - incorporated into the other categories.

I felt that SMART Motors was the least well-articulated category. The US addendum to the report that came out later in 2008 identified four categories. SMART Motors is notable by its omission. The four categories are;

SMART Grid Travel Substitution Road Transport SMART Buildings

I suspect that it is harder for the ICT industry to grapple as fully with the mechanical issues involved in SMART motors compared to the other categories so it is quite possible full justice was not done to this area and Ken's thoughts might go some way towards stating the case there.

I should state that BT is a member of GeSi and the Climate Group and was involved in the global report which I think was an excellent positioning report for the ICT industry. I wasn't involved in it myself though.

Kamal Shehadi, Chairman of Lebanese Regulatory Authority Discusses Forward Looking Goals

Kamal Shedhadi: I have been following with great interest the exchanges and the debates. I just realized that I never posted anything to introduce myself. So here it is:

Kamal Shehadi is the chairman of Lebanon's Telecommunications Regulatory Authority (established April '07). Prior to that, Kamal advised governments, operators and regulators on telecom regulation and strategy and licensing projects in various countries including Jordan, Saudi Arabia, Morocco, Algeria, Egypt, Bahrain, United Arab Emirates, Turkey, and Lebanon. As chairman of the TRA, Kamal's main concern is how to get infrastructure-based competition going in Lebanon in the shortest possible span of time. This requires not only having optimal licensing and regulatory conditions and processes; it also requires having a proactive strategy to pre-empt opposition from a government that is over-dependent on telecom revenues and an incumbent that is - well, like all other incumbents, more interested in "fighting the last war." Kamal and his team have been working actively on ensuring that new entrants have access to existing ducts and conduits and that new broadband networks provide "open access" to any service or application provider.

I apologize for the delay in sending this intro. Once properly introduced to the team, I will jump in and engage in the dialogue.

Harold Feld: If you will forgive my ignorance, does Lebanon have a Commission, like we have in the U.S., or is this a government ministry in which a single minister has regulatory authority?

Shehadi: We do, indeed, have a Commission (called TRA). TRA is governed by a 5-member commission, all full-time, appointed for a 5-year term, non-renewable. Commissioners cannot be removed from office unless they are indicted by a special tribunal (including the three highest judicial posts in the country). The commission is separate and independent from the ministry. However, since the ministry is still the main operator of the PSTN network and, for now at least, the only provider of international bandwidth, and since the ministry (and its public sector entity subcontractor, Ogero) behaves like any other incumbent, this is creating serious friction between the TRA and the minister/ministry. Our job will be made easier when the telco -LibanTelecom- is established.

However, we are using the fact that telco assets are still public assets to push for infrastructure sharing for passive elements and to provide access to existing conduits and ducts to the new entrants to lower the Capex and speed up deployment. It is easier to claim that these assets, since they are still publicly-owned, should be used in a way that benefits consumers, new entrants, etc., and not just the incumbent's market share...

Van der Berg: Good luck to you then. It seems like you have some of the right ideas already :-) Infrastructure sharing is a hot topic here on this list and most people here are proponents of this idea. I hope you can break up the market and establish a way

for multiple telco's/ISP's to make use of the infra and to roll out DSL/VDSL/FTTH in Lebanon.

I hope we can learn from each other.

What are the current hot topics in Lebanon?

Shedhadi: Other than gloating about how well our financial system is doing (which is true, by the way)? How safe and sound it is? Etc.

In telecom: 1. The National Broadband Licenses that I mentioned, which will, most likely, be limited to 3 (including one for the incumbent once corporatized) for the next five years, will have build-out requirements, and open access obligations

2. The spectrum refarming: over the last few years, a handful of private DSPs were able to get all the spectrum they want depending on political maneuvering, including much of the 2.6 and 3.5 bands and all they paid for it was a one annual fee of 75K plus 20% revenue share, thus rewarding spectrum hoarding and limiting competition.

TRA has started a refarming process with the intention of reclaiming much of that spectrum and leaving the incumbents with the bare minimum to continue to operate their networks and expand them. This is the most difficult exercise

3. The privatization of the two mobile networks, initially scheduled for end 2007 and now awaiting the new government that will be formed following the expected parliamentary elections

Feld: Does your jurisdiction extend to spectrum matters as well?

Shedadi: It does. But we have inherited a huge mess that needs sorting out.

Goldstein: Hi Kamel, happy to make your acquaintance!

It seems to me that since there is not yet a private "incumbent", and it's still in the ministry's hands, you have an opportunity to structure the privatization in a manner more befitting of the national interest than many earlier privatizations.

Infrastructure sharing is necessary, of course, because it is a natural monopoly -- requiring competitors to put in their own ducts and conduits is economically not feasible, and harms the streets in the process. But the natural monopoly in many cases (probably to all buildings except the largest) really extends to the entire loop plant as well. So why not define the whole outside plant (conduits, poles, copper and glass loops) as the infrastructure, and put it in a separate entity from the one that provides network services (which gets the switches and existing PSTN customer accounts)? This is sort of how BT structured OpenReach, and a lot of us have been advocating this for the US.

James Seng: Please do not make the same mistake in Singapore during privatization.

Privatize the service but keep the infrastructure. It will save you a lot of headache later.

Verwayen on Net Neutrality

Not What one Would Have Hoped for from the Former CEO of BT

Hendrick Rood Ben Verwaayen releases his ideas on Net Neutrality. Read it, and one can grasp why he in the past was so fiercely opposed to open networks.

Disclaimer: I worked from 1990-1995 for KPN Telecom, where he was then general manager. As more and more parrots of his views started to move up in the rank and file I decided it was time to move on, and joined Stratix. Verwaayen is to me the stellar example of a manager that is able to run a firm that excels in sustaining innovations, improving established product lines, but fails to recognise disruptive innovations. That makes him 95% of the time correct, but he risks a firm he leads to be hit by the 5% of innovations that transform the market radically. Like Lucent and Alcatel missed the router market focusing on ATM, while listening to incumbents.

Rebecca Rijnders: *Ben Verwaayen of Alcatel-Lucent sees the differences between fixed and mobile internet disappear. It will then become the question whether we will pay in the future for consumed datatraffic or for usage of digital services.*

The Dutch CEO of Alcatel-Lucent in an interview with Emerce. said

"I think we will move to a world where the discernability between fixed and mobile internet vanishes. The web needs the internet for access. There is no difference anymore in the way of Internet usage, because the net will adapt to you as an user. Who will pay for what? Are we going to pay for bits and bytes, like we do now or for the value that we assign to a specific message and therefore pay for priority and quality?

The question is how we are going to pay for that. What is the most important service for one, does not need to be for someone else.

I personally prefer e-mail above video. Why should I pay 24 hours per day for broadband? As long as I as a user can decide on that, there will also not be net neutrality. When I travel to Paris, I can choose today whether I go by train, automobile, bus or plane. It is nonsensical that everything has to be evenly expensive. Google off course attempt to block that. I say: empower the customer to decide how he goes from A to B. Compare it to the physical mail, a registered letter is more expensive and nobody thinks that is strange. That should also be done the same in the electronic domain. One will pay extra for speed or security".

The web changes, partially under leadership of Google and Microsoft, from a set of web pages to an environment of applications. E-mail, for instance, shifts from a local computer to the web with services like Windows Live Mail, Gmail en Yahoo Mail leading. Those services are then used, all or not in combination, with online photo books, profile sites or word processors. More and more the PC is not the processor and store of data, but the internet.

Verwaayen poses a critical question: "This so-called cloud computing is a great development, but in various countries the law says that all data that is assembled on you have to remain physically in that country. But where are those data when we are operating cloud computing? And what do we do with se-

curity and privacy? The physical world is still the benchmark for those issues. That security is key.

In Davos a CTO of one of the largest US software firms asserted that the total value of cybercrime is larger than the total value of all drugs crimes in the world. That will thus be a formidable obstacle for cloud computing and other developments. We as an industry do have a great responsibility to take measures with all these new developments to improve privacy and security considerably, because people do trust everything to the net".

A full interview with Ben Verwaayen will appear in Dutch in the forthcoming issue of Emerce Magazine.

Budde: I don't quite follow him Hendrik when he uses his travel example (Paris). Telecoms customers have the choice/or will increasingly get the choice to select high speed and low speed services, according to their needs and their purse. If you only want email get a 256Kbs service for free or perhaps $9.95, full HDTV video will require a higher speed/more expensive package, further more there is wireless, you can go to libraries, Internet cafes, etc, so what is the argument?

This has nothing to do with Google, Microsoft and so on. We - the users - decide what we pay for and there is no need to tax the providers through Net Neutrality.

But perhaps I am missing something?

Goldstein: "Neutrality" is a United States issue. What you're missing is that we do not have a very competitive market here. The typical urban consumer now has a choice of two ISPs, maybe. Three in a few places, one in quite a few. Both charge about $40/month, depending on burst speed and whatever else is in the bundle. There's mobile data too, but it is for the most part ridiculously expensive (when they figure out how to charge), not very fast, and subject to draconian use restrictions.

The Stimulus bill contains a little bit of grant money to finance new broadband builds in unserved areas. That will create ISP monopolies where there is no service now except dial-up (which in some rural areas is a monopoly too). The bill as finally worded has precisely zero requirements for open access (choice of ISP).

Rood: Fred, when I did read the USDA/RUS part of the stimulus bills send later over this list by Jim Baller, I saw the requirement of multiple ISPs as a precondition for getting finance. Page 5 lines 10 to 13.

Goldstein: It's not a precondition. It's a "priority". But there's a whole shopping list of priorities, and it's unlikely that any one proposal will hit them all. Line 17 has a preference for former or current RUS borrowers, who are mostly the current Rural ILECs (subsidy whores). They are unlikely to be open if they don't have to be -- some of the USF subsidies may have required it (the subsidy was for "telecommunications services", not "information services"), but if you give them money without that string, they'd be happy to shut out ISPs.

Earlier **Goldstein**: ATT made some serious noise about charging Google and others who are *NOT* their customers, but who reach them via peering, for "access" to their subscribers. Pay your ISP, pay your customers' ISP (ATT) too. They've also been active in the IMS world, where the design is for a metered-by-"value" application layer relay. This happened right after the FCC, at the ILECs' request, ended broadband common carriage, so customers lost the right to choose what ISP they wanted to reach over ATT's wire. This is what led to the neutrality wars.

Once "neutrality" became a political issue, companies figured out that it could be used as a wedge against standard ISP Terms of Service. After

all, it's not "neutral" to ban file servers. So Vuze took the lead and got permission to build a content distribution network using residential broadband subscribers' PCs. Now CNN distributes some popular video feeds that way. It's grossly inefficient but the ISP pays, not them. And while ISPs are responding with gigabyte/month limits, the ortho-neuts are crying in their imported beer that it's somehow wrong.

Rather than craft a set of consumer-friendly solutions that might actually protect ISPs against such things, it has become a battle of extremes. Verwaayen seems to be firmly in the IMS camp, wanting to replace the Internet (a business model, not a protocol suite!) with IMS, an equipment-vendor-friendly approach where he and his carrier-customers can sell individual applications. Just like 1982's Compuserve, but with video.

Rood: I think this is a crisp summary of Verwaayen's mindset.

He sees the Internet business model as Fred describes it as flawed, as he cannot imagine any end user desiring the freedom in choice of ISP as a business model, because he is thinking about end users in terms of only desiring applications and content provisioning and then the IMS model is the only one that makes sense to him.

Goldstein: Indeed. In his mindset, he gives the users a choice between one or two "news" sources, an email-like service, and a few games and TV shows, and figures he's given them all the choice they need. It's 1980s Videotex (Minitel, etc.) writ large.

February 15 **Kelly**: Verwaayen was CEO at BT for 5 years of my tenure.

He was the driving force behind the commitment and openness that saw the creation of Openreach and equality of access to all operators, new and old.

He supported the approach to open networks and systems that pervades BT's 21CN.

And he lobbied across the world for openness and competitive access to networks in the same way the UK is moving towards.

Before you jump in to criticise him (we all have facets of our business lives that others can argue with), understand what he's tried to do to further the goals of openness and choice in the market.

Goldstein: I wonder how much the culture makes the man, or the man makes the culture. Clearly Verwaayen today is spouting the same party line that inculcates Lucent culture. Lucent spun out of AT&T, retaining much of the old monopoly (pre-1984) culture. It gave lots of vendor financing to damn-fool CLEC projects in the late 1990s, got burned, and largely wrote off the CLEC market, focusing on wireless infrastruture and incumbent optics. I doubt Alcatel's heritage is that much different, though it was never owned by the PTTs that were its masters. At AlcaLu, he's naturally going for the most equipment-intensive solutions. And those certainly aren't the most open.

Maybe at BT he wasn't actually so much driving openness as recognizing that it was a necessary condition for surviving the inland market.

I think of Sam Simon, ILEC PR flack extraordinaire, King of Astroturf. He started off 25+ years ago as a Nader Raider, and started TRAC as a real consumer advocate group. But then he created a business, monetized his credibility (basically turned TRAC into a Verizon front, attacking the other LD carriers), and became the monopolist's best friend. Sometimes a really fat paycheck can change someone's perspective.

Rood to Joe Kelly of BT: Joe,

Alongside his tenure as BT CEO he occasionally gave his views on efforts on the Dutch

telecoms market. He strongly opposed efforts undertaken by municipalities and housing corporations to create open networks at the physical medium layer (structural separation style).

Coincidentally the regulatory affairs officer of BT Netherlands was also a member of the Amsterdam municipal council (Christian Democrat party) and ended up in a rather schizophrenic loyalty position when the vote for the Amsterdam Citynet plan came up in the council.

The openness that BT promotes is an open access to layer 2 access networks.

As far as I can see, BT is striving to create a regulatory model that copies the regulatory model created in the British market to other European countries. But then with the national incumbents in a similar role, however under heavy regulation.

It has fiercely lobbied our regulator during the past years, together with Verizon Business, to start regulating the business market. This despite the fact that BT owns one of the two most extensive non-incumbent fiber networks in the Netherlands to a large number of Dutch local exchanges. It is leasing out this fiber to Telecom Italia subsidiary Bbned (a wholesale DSL CLEC) and uses it for some of it's own corporate customers to create on-net DSL.

With respect to the VDSL push seen in most countries Europe it looks to me that BT has concluded that they will retract their position with ADSL/SDSL networks and rely on wholesale inputs by incumbents. That position effectively re-establishes equipment/wholesale monopolies in most European countries, boosts regulatory interventions and trashes the business of wholesale CLECs.

I am completely aware that this might be optimal for BT's balancing act between it's position in it's home market and it's interests with BT Global Services in the continental business market.

One might however question, whether that position is in the interest of the pan-European business community, who tend to prefer networks provided by a single form controlling the active equipment layer to enable competition between operators in tailored service level agreements.

At the November 2008 ECTA conference it was several times remarked that the passive layer of fibre optic networks can be architected in a pro-competitive design, as well as an anti-competitive design.

As far as I can see, BT has been quite mum about pro-competitive designs at OSI layer 1, but pushes for open access at layer 2. This is obviously that will not exploit the much larger reach that fiber optic networks offer.

I cannot otherwise judge BT's current position as a choice to reduce future network CAPEX outside the UK, while taking an ongoing regulatory hassle to get layer 2 access per country for granted as well as accepting ongoing diverging services specifications.

After my ECTA panel session with Grant Forsyth, I asked him why BT was not actively lobbying ECTA members to hammer out a set of uniform layer 2 wholesale specifications at a pan-European scale and then bounce them with INTUG. He said BT did not do that because there was no regulator at the EU level, contrary to the UK.

I personally think that regulatory vacuum is an opportunity, as the EC would love to confront all European NRAs with a set of specifications (an elaboration of the Metropolitan Ethernet Forum's work) agreed upon between operators and INTUG to avoid the national fragmentation that is now appearing due to the strong focus of NRAs at national consumer market solutions.

Leaving the decisions about how to implement wholesale access at layer 2 to the NRAs creates national regulatory capture of NRAs in hammering out wholesale open access at the technical level, as at country level it will be the local incumbents that call the shots and regulators will go along with them, to the detriment of pan-European cross border businesses.

As not all European regulators are like Ofcom, willing to engage heavily in the process of technical hammering out specs, this creates a multi-year hassle.

The fact that BT is mainly playing the regulatory ball in most EU countries on wholesale access, instead of either pushing for pan-European pro-competitive designs of fibre optic outside plants to enable an active network roll out to most regional centers over Europe in the forthcoming decade, or creating a more uniform set of pan-European agreed upon wholesale access standards, suggests to me that there is still a lot of regulatory addiction and open access spin doctoring. Home markets still prevail thinking.

That might serve European incumbent operators well in the short run in each of their markets and on their balance sheets, but it is a long-run negative for the further development of the EU common market and the competitiveness of European businesses as it will lack operators at (sub-)continental scale that can change and upgrade network services technology without need of asking for permission to innovate to both the local incumbent as well as the then unavoidable NRAs.

About a decade ago, I saw a webcast of the Gilder Telecosm conference, where David Isenberg spoke about his experience after writing Rise of the Stupid Network. As he explained the real killer of innovation in the "intelligent network" is the fact that every idea to innovate required working all decision layers in the corporate hierarchy. A surefire way to kill off most innovations.

Working the regulators in their three-year cycle to get incumbents all around Europe adapting to new (wholesale) services, is adding another set of layers of decision and indirection. It is a boon for telecoms lawyers and regulatory consultants, but it is a burden for innovation and market dynamics.

Although I do earn part of my living with regulatory consulting, I prefer to operate in more dynamic pan-European markets where most time can be spent on developing new products and services instead of regulatory trench warfare.

Up to now, I do not see advancing open access to layer 2 networks as reducing that kind of red tape. Open access to pro-competitively designed layer 1 (fibre) networks instead do offer that kind of opportunity of cutting red tape for advancing network capabilities and specialising services and I have seen Ben Verwaayen and BT opposing it, even when they are in a CLEC position.

I think that amounts to short-term optimising the bottom line and is more a tactical than a strategic choice.

Pimentel to Joe Kelly: Good point and respect is acknowledged.

Apparently now, he is clearly on the side of IMS. On a technical level it will never work as architected. On a business level it favors the monopolistic position of the TelCOs and CableCOs. It shows a narrow view of consumer behavior and open market trend.

The question remains whether the current winds of change will bring "the" perfect storm where open networks and innovators can actively compete with the TelCOs/CableCOs.

Executive Summary

Harvey Newman & LHC, pp 1-30

The April 2009 issue features the first part of a long two-part interview with a high energy physicist Harvey Newman. Harvey explains the origins of his involvement in data networks and other fields of information technology, along with some highlights of his career at the frontiers of modern science. His founding of international networking for his scientific field started with the need to be working with European particle colliders, first in Germany and then at CERN while raising a family, teaching and running a research group at Caltech in Pasadena.

Harvey explains how during the 1980s and then into the 1990s, he worked with the early data network protocols and beginning with X.25 adopting varied protocols like DECnet SNA in TCP/IP as the 80s turned into the 90s -- eventually switching to a reliance on TCP/IP during the mid-1990s. The data output of large particle accelerators continually pressed against the state of the art of information technology and networks. Deregulation in Europe in 1998 made it possible for Harvey to move into the optical hierarchy for his transatlantic links. With the floodgates open and the data needs of the large hadron collider truly unprecedented, Harvey began to build a team that was able to take on the demands of what has become a global collaboration.

He has focused on innovative problem solving in several network-related technologies, developed a new class of global-scale systems to support the science, and delved into areas of science policy as well as telecommunications policy. As an example of this we may point to his design work on the data dissemination structure for the LHC. In this case he has focused not only on issues of technology, but also on how to structure and operate his data dissemination architecture so as to encourage cooperation amongst universities and laboratories throughout the world, while also working in many countries to bring along those science groups with less-advanced national infrastructures, and economies, towards equality with their peers.

A critical skill is his ability to conceive, fashion and implement an architecture that enables scientists located at hundreds of sites in many nations to successfully carry out their work in a decentralized, cooperative way, while remaining within the orbits both of their local universities, all the while connected via networks to their experiments at the LHC and especiallly their colleagues working there are the CERN laboratory in Switzerland and France. The decentralization of the architecture created an attractive feedback loop that enabled the widely decentralized communities, and their funding agencies, to participate in a project of global importance while building local and national capabilities.

Another outcome made possible by carrier competition was the ability to work independently of carrier technology and protocols. Harvey made a critical point when he said: "What has made the wide distribution of data to and from many switches and also many servers at one and now 10 Gbps feasible was the abandonment of the old "carrier class routers" and the adoption of 10 gigabit Ethernet (10GE), in full-scale switches where a 10GE port costs on the order $ 10k." Being able to gain access to basic fiber infrastructure in the form of lightwaves enabled Harvey's global science collaboration to function relatively free of telco business model constraints.

No Progress Until the Telco Model Becomes Technology Driven

Newman: "When it comes to wide area networks, the world is mostly politically and policy driven, and not technology driven; unlike some other areas of information technology. Somebody mentioned today on your list that the carriers seem to be waiting to see public demand. However this is not true for storage and not true for processors in that the makers of these devices just keep increasing their capacity, challenging users to create and develop new applications and new skills. **Unfortunately the deployment of bandwidth is not technology driven. The vendors are looking at markets.**

We seem to have a kind of conundrum with bandwidth: you do not seem to have an elastic demand. If the producers put out 10 times the capacity, the public seems to be unaware that it's there and unsure of exactly what they should do with it. With bandwidth neither awareness nor a real cost model is driving an elastic demand."

Indeed the telecommunications driven market gives us billing for every bit, carrier class routers at $ 1.5 million a pop and nothing more innovative than triple play. Harvey's world unfortunately is a much too well kept secret. As we shall see in part two next month he is using the needs of his global science collaboration to inspire, or press, political authorities to be innovative in building new infrastructure. As we cling to the last vestiges of our market model other countries – less restrictive – will reap the benefits on a large scale - not only in science, but also in education technology and economic competitiveness.

Harvey Newman: It's a different strategy **based on dark fiber infrastructures which have been deployed by national research and education networks (NRENs) in a growing number of countries outside the US, including economies like those in central Europe (Czech Republic, Slovakia and Romania for example) that are orders of magnitude smaller than the US. Just put things out there. And find out what people want to use. Create conditions where the costs are low so that mass-produced technologies are able to use the infrastructure fully, and the methods to use it effectively are created rapidly on a large scale, and evolve as the infrastructure evolves.**

Cook Report: the problem is this won't happen as long as national backbones are powered by ridiculously expensive carrier class routers supported by customer billing systems designed for the days of analog technology.

The final half of the first part of the interview with Harvey explains the technical workings of the collider itself and lays the groundwork for understanding the LHC network and its unique accompanying telepresence system that will be the subject of next months conclusion.

IPv4 Market Risks, p. 33

As with new registry policies IPv4 blocks become property, the question is raised as to whether there will emerge a business model focusing on the aggregation of owned IP blocks into routable chunks. **David Conrad** responded The fact that speculators will have locked up the allocated-but-unused address space merely maintains the status quo...

Vest: Actually, the status quo is likely to change in one huge way, with multiple far-reaching effects, regardless of whether the market "works" and liberates lots of additional IPv4, or absolutely fails to motivate current "reserve address holders" (i.e., hoarders), or is fatally distorted by speculators, etc.

The next time a lawsuit arises over control or use of IP address resources (it has happened several times in the past), at least one side is going to have recourse to the protections afforded by whatever property law is in force in the relevant jurisdiction. Before transfers, resource users were not considered to "possess" IP number resources in the way that makes the phrase "possession is nine-tenths of the law" such a perennial favorite. Once transfers start happening, and lawyers and accountants are forced to reckon with the implications for both transfer transaction participants and everybody else, this will no longer be true.

The first likely impact will be to render moot any policy-based requirements or restrictions on IP address transfers themselves. The really big impact will come when this starts to undermine participation in (any) shared public registration database -- which to date is the only thing that has preserved the presumption of uniqueness that puts the "public" in "public IP addresses". If that goes, the only things that might be able to put it back will be property law and/or public regulation -- which will have to be reconciled with other laws about privacy, cross-border trade and direct investment, etc.

Editor: What is being abandoned according to **Tom Vest** is "That canonical RIR-era arrangement -- hierarchical routing and prefix-length filtering for the big operators plus a neutral, open allocation mechanism for new entrants -- created an excellent, conflict-minimizing industrial environment for growing the Internet." snip "Now, however, that bargain has been abandoned -- and it seems highly likely that the party that lost out (i.e., the new entrants) will ultimately react the same way that they do in other industries that are shaped by a concentration of market power/control over critical bottleneck inputs.

So, in the end, we may get to run the market advocates' experiment and see exactly what the routing table looks like when there are no durable limits on demands for finite routing system capacity."

And later **Vest**: **The demand for confidentiality or non-transparency, and for freedom from "bureaucracy" is common if not universal, and quite understandable in most cases. Whenever transacting parties can get away with non-disclosure, they usually do, for better or worse (e.g., in the banking sector, better yields right to individual privacy, worse yields global financial collapse). In this particular case, the complete absence of any countervailing enforcement mechanisms is sure to encourage more people to press that demand. But as a result, the registration database that is the only mechanism that assures the uniqueness of IP number resources (and hence their basic utility) will cease to be sustainable.**

KPN Joins Amsterdam in Extending its Fiber Build p. 43

Paul Budde on Feb 4*: The City of Amsterdam announcement to now move to the next stage of their FttH project - with another roll out covering 100,000 connections - is a clear indication that the concept of open access FttH networks is a valid one. This will have large-scale implications for countries around the world that are looking at using open network based telecoms infrastructure projects to stimulate their economies.*

BuddeComm has been involved in industry policy discussions with experts on three continents about the future of telecoms. We have developed a scenario. The

vision naturally has a range of strategies attached, but in simple terms it works as follows: **1.** *Telecoms infrastructure is of national economic and social importance;* **2.** *For our societies to profit from the digital economy infrastructure must be based on the principle of open networks;* **3.** *This allows us to multiply the benefits this infrastructure has to offer to other sectors such as healthcare, education, energy, environment, media and communications;* **4.** *Once open networks and the access tariffs are established the national telecoms operator will be in the best position to run this network, thus avoiding the necessity for wasteful duplication.* **5.** *Infrastructure and digital applications will need to be developed parallel with each other and this requires trans-sectoral thinking from the government and the industry - not the current silo thinking.*

While it is great to put such visions in front of the policymakers in the end the question is whether this is achievable - can such a vision be implemented? Again the Netherlands is paving the way. Amsterdam was one of the first to identify the multiplier effect and the need for trans-sectoral thinking. Unfortunately the incumbent KPN first preferred to roll out fibre nationally **[Snip]**

Editor: The Open network should become a magnet for all providers and the supplier that keeps it open gets the maximum return on the cost of the build by attracting many new businesses. Think in terms of the shopping mall that grows because of all the other places that one can do business with on a one stop basis.

Vincent Deker: **If you allow all your competitors on your network, all services will come on your net, and that results in the lowest cost possible per service. Which in turn attracts more customers for those services, so your network grows much faster. An open network is not charity from our side, in the long run it simply works best for everybody.'**

IT, Internet and the Power Grid p. 47

A discussion about the use of the internet for energy saving and CO2 reduction shows the complexity of needed planning.

Bill St Arnaud: Most people think today's Smart Grids will reduce consumption- where in fact they largely only displace consumption from peak demand. In the several studies done s far on smart grids and smart meters, the savings have been very small for consumers.

Some early studies indicate that providing feedback to consumers on energy consumption will reduce demand. But these studies are very limited scope and the sample population is usually made of dedicated and concerned homeowners. We will need longer-term studies with larger populations to see if feedback meters will have a meaningful impact.

The biggest beneficiary of smart grids is not the consumer, but the utility who does not have to build more power plants for peak load. Indirectly this marginally does reduce CO2 emissions because base load is usually nuclear while peak load is usually gas powered plants

Jeff Sterling offers an excellent summary of three different ways of looking at he concept of a smart Grid.

Smart Grid is a Vision. As we speak, smart grid doesn't really exist and it's probably an evolutionary process. The question is whether smart grid is a transformative conversation with transectoral implications that goes beyond the boundaries of the IEEE or other private industry group that are looking to optimize the electric utility system. I would argue that smart grid is the connective tissue we need to enable the new energy sector.

Smart Grid is an operating system. The bottom line for any successful implementation of smart grid is software based. There are firms out there that want to be the Microsoft of the Grid. Google, Cisco, IBM, Microsoft, and a host of startups are all trying to figure out what to do. What is the model? I would hope for a open standards approach that would create a development platform for localized energy projects.

Smart Grid is a community network. Smart grid is the ultimate opportunity for transectoral integration at a community level. Unless you include the human factor into the system you are reinforcing the silo mentality that is ossifying innovation in community infrastructure.

Progresive Lebanon Regulatory Agenda p.58

Kamal Shehadi is the chairman of Lebanon's Telecommunications Regulatory Authority (established April '07).

Shehadi: We are using the fact that telco assets are still public assets to push for infrastructure sharing for passive elements and to provide access to existing conduits and ducts to the new entrants to lower the Capex and speed up deployment. It is easier to claim that these assets, since they are still publicly-owned, should be used in a way that benefits consumers, new entrants, etc., and not just the incumbent's market share...

Van der Berg: Good luck to you then. It seems like you have some of the right ideas already.

Goldstein: It seems to me that since there is not yet a private "incumbent", and it's still in the ministry's hands, you have an opportunity to structure the privatization in a manner more befitting of the national interest than many earlier privatizations.

Infrastructure sharing is necessary, of course, because it is a natural monopoly -- requiring competitors to put in their own ducts and conduits is economically not feasible, and harms the streets in the process. But the natural monopoly in many cases (probably to all buildings except the largest) really extends to the entire loop plant as well. So why not define the whole outside plant (conduits, poles, copper and glass loops) as the infrastructure, and put it in a separate entity from the one that provides network services (which gets the switches and existing PSTN customer accounts)? This is sort of how BT structured OpenReach, and a lot of us have been advocating this for the US.

James Seng: Please do not make the same mistake in Singapore during privatization. Privatize the service but keep the infrastructure. It will save you a lot of headache later.

Verwayen Takes Pro Incumbent View of Network Neutrality at Alcatel Lucent p. 60

Goldstein: Rather than craft a set of consumer-friendly solutions that might actually protect ISPs against such things, it has become a battle of extremes. Verwaayen seems to be firmly in the IMS camp, wanting to replace the Internet (a business model, not a protocol suite!) with IMS, an equipment-vendor-friendly approach where he and his carrier-customers can sell individual applications. Just like 1982's Compuserve, but with video.

Rood: I think this is a crisp summary of Verwaayen's mindset. **He sees the Internet business model as Fred describes it as flawed, as he cannot imagine any end user desiring the freedom in choice of ISP as a business model, because he is thinking about end users in terms of only desiring applications and content provisioning and then the IMS model is the only one that makes sense to him.**

Contents

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

This issue leads off with the first part of a two part interview with Caltech physicist Harvey Newman and four weeks of symposium discussion - including more discussion of the policy problems inherent in emerging IPv4 adress markets.

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 May 2009 issue - **out about March 28 the second part of the interview with Harvey Newman. I anticipate the delaid nterview with Frank Coluccio for the June issue.**

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