Research on Networks versus Networks for Research: The Need for International, Internet, Testbeds

Peter T. Kirstein
Department of Computer Science
University College London

ABSTRACT

The paper reviews the early history of Internet testbeds, and indicates where they have played an important role in advancing Internet technology and usability. The current status of national testbeds is discussed, and the lack of an international equivalent is pointed out. Finally, we propose a medium scale, advanced testbed, which would include several types of technology (including cable, satellite and mobile), include several carriers and countries, and allow experimentation with the latest technology that could break the network.

1. Introduction

This week, the University of California in Los Angeles (UCLA) is celebrating the thirtieth anniversary of the first four nodes of the Arpanet. Over the past thirty years, the Internet has grown from 4 nodes to several hundred million; this represents a compound growth rate of over 80% p.a. every year! The equivalent growth in traffic is nearly doubling each year over the same period. This compares with a growth of the number of the telephone network of around 16% p.a. over the last 120 years. Clearly, the main advances in the network technology have come from the US. This has been greatly helped by the existence throughout this period of major networks that could be used as playthings or testbeds for experimenting with new techniques, technologies and applications. These testbeds have not been uniformly useful; that is the wrong thing to ask. If each testbed had been useful, then they would not have been ambitious enough. However even many of those that appeared not to have been successful at the time, have often led to further massive improvement in facilities in the future. It is one tenet of my talk that such testbeds are invaluable drivers of the technologies; their paucity, or at least their more limited nature, elsewhere has contributed to the relative immaturity of the activity in many other countries.

Many lessons can be drawn from the study of the Internet. Here I will concentrate on four: Openness, Heterogeneity, Availability and Globalisation. I will give examples of how each has benefited the system that we now call the Internet, and draw some lessons on how they may continue to do so. Finally, I propose an initiative on a new, medium-scale, very heterogeneous, international network and application testbed.

Because I am more concerned at the concept of what can be learned from testbeds than the testbeds themselves, I will give only limited references. A much fuller set of references on the early experiences is given in [kirs].

2. A Potted Summary of US Internet Testbeds

There has been large-scale network activity both by industry and different segments of the US government since the sixties. Nevertheless, it is recognised universally that DARPA has been the main driving force. It is no coincidence that of the ten previous recipients of the prestigious SIGCOM award, three have been in DARPA at some time, and several more have been associated closely with Arpanet or the Internet. DARPA supported a significant number of large-scale data network testbeds, where *large* must clearly be considered relative to the then size of data network. These included Arpanet (68-79), Packet Satellite (75-80), Packet Radio (73-78), Wideband satellite (78-88), Internet (77-87), Gigabit (91-96), DARTNET (91-96), CAIRN (96-99), and Glomo (93-98). All of these started off as technology demonstrators; many were encouraged to develop into more generalised application platforms.

In the early 70s, many of us were using the fledgling Arpanet for long terminal sessions and even packet speech. These were not done as demonstrators; they were for serious work outside the Computer Science Area. This was rapidly replaced by the electronic mail systems of the middle 70s. By 1978, there were already a hundred UK research groups approved for regular collaboration with partners in the US. The simultaneous need to develop the technology to build a high quality network, and to provide high quality services on a significant scale were so important. They led both to the Arpanet becoming such an excellent basis for the Internet, and to the early Internet being extensible into its current commercial success. To give just one example, I remember an early meeting around 1972 where the Arpanet had had problems, and ISI had been unavailable for a week. Since all the DARPA Management Information System was at ISI, this was an intolerable disturbance. Larry Roberts said to the group "You will stop this happening again, or I will withdraw all support from Arpanet". As a result there was a massive attack on all the areas that could cause significant problems, and a step-function improvement in availability. It was this drive to provide leading edge technology, while exercising the system, which was so important in these testbeds. All these networks were *research networks* and yet they migrated also into *networks for research*.

There was an important hiatus in the mid-80s. DARPA felt it had done the basic research on networking technology; there was little need for *research networks*. The provision for *networks for supporting research* was considered the remit of NSF for the academic community. At the same time other Federal agencies were putting in networks both to support research and even operations; these included Energy, Defence and NASA. The scale of these operations provided applications drivers, even if the research networks had these only on a limited scale.

The testbeds supported by DARPA in the 90s have had less of an applications driver. Many of the Gigabit testbeds did have one or two target applications; however it was possible to tailor many aspects of the testbed to the application, and few large-scale persistent applications emerged. The requisite network components were developed, and the basis was set for another push into advanced applications. The ATDNet and ACTS networks did incorporate applications; however they were much less available to the bulk of the research community.

With the formation of the Federal Network Council (FNC) in 1995, followed by the Large Scale Networking Group (LSN) in 1997, the US seems really to have put its programme in place. With the different phases of the Next Generation Internet (NGI) programme [ngi], there are a number of research networks. With the rapid expansion of the VBNS network, followed by Abilene and the Internet-2 [int2] initiative, there are networks to support research. These have much less advanced technology, relative to the current state of the art; it follows the now well-established US successful formula of having sizeable advanced technology research networks, with their technology feeding through into yet large *networks for supporting research* with relatively stable technology. Moreover, Internet-2 also has a significant number of advanced development activities like IPv6, Measurement, Multicast, Network Measurement, Network Storage, Quality of Service, Routing, Security and Topology.

When we examine these testbeds by the criteria of Section 1, they do very well. Almost all were very open; being funded partially by US Federal funds, the results were normally made available widely. Their heterogeneity was limited. While many sorts of computers were often involved, the assumptions were often very similar. This was partly due to the national character of many of the testbeds, and partly on the way the funding was provided. Availability was normally good; of course technical faults might limit availability, but most were designated to be available most of the time. Globalisation was very limited; the experimental networks were mainly restricted to the US, with certain notable exceptions - some of which are discussed in Section 4.

3. The Value of the US Testbeds

Most of the US Internet testbeds had three characteristics:

- **Broad Coverage** For the area where the testbed was being directed, many of the interested parties could play.
- **Openness** Most of the Internet testbeds were Government financed, so that the results and even implementations were publicly available. While this did not apply to big networks like BITNET/EARN or DECNET, these were not really intended as testbeds. They were large-scale networks with the main provision of the software from the supplier. There were some very notable exceptions to Openness; many of them sparked off interesting technical developments (see below).
- **National Scope** With some very limited exceptions like SATNET, most of the Internet testbeds were national. When there was an international dimension it was very limited. This was partly due to the costs of international communications, and partly because the national funding made it inevitable that most of the development would be national.

When the US testbeds had an international component, it added an unexpected ingredient to the project, which led significant new elements.

4. Variations introduced by the International Dimension

4.1. International Aspects of the Early Arpanet

In 1973, when there were about 20 nodes on the Arpanet, two additional ones were added in the Norwegian University Computing Centre (NUCC, Kjeller) and at University College London (UCL). Although both were equipped with the same hardware, an Arpanet Terminal IMP (TIP), the two developed very differently. The NUCC connected in two computers for the NORSAR Seismic Array, and contributed significantly to the internationalisation of seismic monitoring. However, its TIP was not connected to the public telephone network, and academics could only use it in-house. It never made any new demands on the Arpanet testbed, and did not participate in any research (except SATNET, see below). The UCL node was connected to the public telephone network, connected in a large computer sited 50 miles from UCL with non-Arpanet protocols, and put in access control at a very early stage. It provided services that grew from a few dozen in 1973 to several hundred research groups five years later. I do not want to overestimate the impact of this small perturbation on the Arpanet Project, but amongst the novelties it introduced were the following:

- The first siting of a large scale remote system with protocol translation of terminal and file transfer in a gateway;
- The first application of access control both for logging on to a TIP and later for international services;
- The first use of traffic aggregation to reduce bandwidth demands. This included making the character at a time operation of the time-sharing PDP-10s work in a line-at-a time mode, and forcing mail-bagging of outgoing mail with distribution expansion nearer the destination;
- One of the first linking together of whole networks (Arpanet to the UK SERCNET)
- Probably the first systematic translation between Host Directories (the Internet Domain Name System to the UK Network Registration System).

Each of the above uncovered needs that would have surfaced sooner or later on the national scene. However, the heterogeneity introduced by the international dimension, without the rivalry brought in by a competing local activity, as between early US agencies, encouraged the enrichment of the activities on both sides of the Atlantic. Clearly the British gained greatly from this; equally clearly the positive nature of the interaction arose largely from the attitude of our US partners in wanting to help our activities.

4.2. SATNET

Another interesting testbed was SATNET. This testbed was essentially international, and raised many new issues directly because of this. This testbed was based on Satellite IMPs being sited in the satellite earthstations of the carriers belonging to British Telecom, COMSAT, the Scandinavian satellite carrier, Italian Telecom and German Telecom (actually DFVLR). The technology required that the equipment actually be sited on the carriers' premises, in the direct chain to the up-link. It was new to all these carriers to have equipment owned by a customer and remotely managed on their premises. The importance of this testbed was only partly the technology it introduced. The way that the satellite bandwidth was managed between the earth segments was used a lot later. More important was the way that gateways were introduced between the Arpanet and the SATNET segments. The reason that Bob Kahn wanted the gateways was to de-couple the development of the facilities concerned with the satellite system from those concerned with the Arpanet. It was also to open up the development of the network technology to others than Bolt, Beranek and Newman (BBN) who had developed the IMP, Terminal IMP and Satellite IMP. While in fact BBN still provided, from a separate group, this gateway, it did provide a separate component that could be opened up. Indeed, David Mills produced his Fuzzball to rival the BBN gateway, which was used also in the early stages of the NSF networks. That separate component later became the router, which allowed the whole new industry to be formed, and is largely responsible for the existence of the Internet today.

The existence of the SATNET nodes on their premises brought the Arpanet much more directly to the attention of the different carriers. The Germans were not interested at all, and there was no national access to the German node; such access had to wait until CSNET came along from NSF nearly ten years later. The Italian interest was weak; the satellite carrier was a different company from the national one. No real research developed from this work. Even the Scandinavian interest was not too strong; because of the lack of broader access in Norway, there was no strong user pressure. By contrast, British Telecom became very interested in the potential of the networks, and of their connection with their own packet-switched networks. For some five years after the experimental phase officially ended, the SATNET path was an alternative to a use of the British Telecom international packet-switched service.

Another aspect of that early period was the need for the carriers to face the tariff impacts of the digital services. While voice was normally analogue, most modems operated at 9.6 Kbps with a few complex modems at 48 Kbps. The SATNET service was one of the first users of pure digital services, and used a multi-destination, half-duplex service. Because it was digital, all voice users were provided with 64 Kbps service - though 64 Kbps digital service on analogue lines was being charged at 4 - 10 times the voice rate. In the early days SATNET was regarded as only an experiment, and so the tariff concerns could be overlooked. Nevertheless, these concerns were brought early to the attention of the carriers, so that the whole question of shared usage became easier when a larger scale international Internet grew in the 80s.

5. Differences in US, Canadian and European Research Net Support

The early support for Arpanet and Internet came almost entirely from US Federal funds. It included some industrial laboratories (e.g. Xerox Park and IBM Research) at their own expense; almost all industrial participation was under Federal contract. Even the early communications costs were largely at commercial rates. It was only with the Gigabit testbeds that there was a large-scale investment by US industry in cost-sharing arrangements. With their large investment in fibre, and a considerable surplus capacity on many routes, the US carriers have continued to participate on a shared cost basis - but almost only domestically up to now.

In Canada, there was an attempt in the early 70s to set up a Canadian Research Network (CANUNET), under the auspices of the Department of Communications. This had the political aim of unifying Canada at that stage of its history. For various reasons CANUNET was never built. It was only in the 90s that Industries Canada was given the responsibility for constructing CANARIE [can]. This is a true testbed, and has substantial contribution from Canadian industrial organisations. Moreover, it is available both to the Canadian universities and their industries. It has truly experimental components, while at the same time providing a proving ground for advanced applications.

6. The European National Testbeds

The Europeans research community has almost always had a more unified view of networks than the US. There has normally been only a single agency responsible for networks for research – and that has almost always been responsible also for research networks. There were some early research networks; one of the earliest European ones was the European Informatics Network (EIN). Most of these early network were either *research networks,* or *networks to support research.* It is very rare that they have been both. Thus in France the Cyclades Network was a research network; neither EIN nor Cyclades had serious users. By contrast, Euronet was based entirely on the Cyclades technology, and was put in to support researchers; it had no element in developing network technology.

Almost all the academic network put in during the 80s and early 90s were *networks to support research.* In general they had a very small ostensible research content. The British SuperJANET was supposed to be an exception. Its original concept was that a clear portion of the network would be made available as a *research network* for network researchers; it was on this basis that British Telecom provided the special rate for the network, and agreed to fund academic research using it. In the event, certain high visibility applications, in particular video conferencing, were found to be very sensitive to traffic variability; as a result, none of the network was ever really made available as a research network that could be broken.

Since 1995, two sets of networks have been made available as real testbeds. One set comprises certain national research networks - in particular in Scandinavia, portions of SURFnet in the Netherlands, and some research networks in Germany. Clearly CANARIE [can] in Canada has had this mission, and there have been some bilateral arrangements between German Telecom and Teleglobe to connect the German DFN [dfn] and CANARIE for specific experiments.

The European Union intended to provide a *network to support research* for its research projects under the Framework IV programme. The resulting JAMES network fulfilled this function for several years; it also allowed some very limited network research - particularly on ATM interconnection. However, because of the way the constituent components were run by competing carriers, it was difficult to get much real insight into the operation of the network, or to have outsiders become much involved in experimenting with the network itself. The successor to JAMES, the QUANTUM project [quan], is again principally a set of interconnections between the national service networks; it will allow limited experimentation outside the network, or on Virtual Private Network (VPN) components that can be derived. However these are of very limited scope; they are similar, but much less ambitious, than those envisaged on Internet-2 without the added research potential offered by the NGI testbeds [ngi].

With the growth of the Internet, the carriers and some national academic network providers have started to see the need for serious research networks.

Several carriers have started activities for research networks. UNINETT/Telenor (Norway), SUNET/Telia (Sweden), DFN/German Telecom (Germany), British Telecom (the UK) and SURFnet/KPN (Netherlands) have significant projects. In most cases, the projects are very limited in both size and scope. They normally do separate the applications from the network technology; only very limited scale network topologies are available for experimentation. All are restricted to national activities - though Nordunet, DFN, SURFnet and SuperJanet have both joined Internet-2 and planned connections to it.

I separate out the British component in the above from the others for a specific reason. Most of the national activities are a combination of a carrier and the organisation responsible for networks to support researchers. The early SuperJANET was just such a collaboration. However, when it was found not to incorporate real networking research, British Telecom decided to pursue that avenue by a separate network in conjunction with a few particular universities. A Proposal for a world-class testbed are currently under consideration in a way that would include both UKERNA and private industry.

Some of these networks are experimenting with advanced access arrangements - in particular xDSL and even wireless access. Only the French and some Germans seem to have an active programme with DBS satellite; in both cases these seem to be connected to some political considerations. In the French case it is connected to links with the Francophone countries, and in Germany it is connected to links with the former CIS countries.

7. Industrial Collaboration in Europe

In Europe there has been a massive push, under the European Framework programmes, to encourage collaboration on a pan-European scale between partners in different countries both industrial and academic. In theory, the JAMES and QUANTUM networks were intended to make on-line collaboration easy in such projects. In practice, this has not yet happened. The access to JAMES and QUANTUM has been mainly through the academic networks; they have not necessarily provided either easy or low-cost communication to commercial entities. JAMES was not organised in a way that such access could be guaranteed either. QUANTUM is just an interim project until the FRAMEWORK V initiative gets properly underway; whether it will provide the right environment for the participation of the commercials is not yet clear.

So far most of the collaborations have made only very limited use of on-line networks. Even in the ACTS programme, which was designed to develop a Broadband Integrated Network Service for Europe, most of the final demonstrators were built up at single sites!

There is an additional problem with the involvement of commercial companies in the FRAMEWORK projects. Most of their research activities are only 50% funded by public funds, and the results are not necessarily published widely. Easy dissemination of software cannot be taken for granted nearly as much in these projects as in their US equivalents.

In theory, the ACTS programme envisaged the construction of National Hosts. A National Host would consist of some specific network-based research, and it was envisaged that the National Hosts would be connected together. Both TEN-34 and JAMES were supposed to connect together these National Hosts. The theory was excellent; if completely carried out, it would have rivalled the plans in the US. The national academic networks were mostly offered as National Hosts. Other important facilities like national cellular telephone testbeds, xDSL testbeds and experimental satellite facilities were offered also. Unfortunately, the reality was far different from the 1994 hopes. Most of the interesting experimental resources were not linked in to the national academic networks. Moreover, the lack of a real pan-European research network did not allow these interesting resources to be linked on a trans-European basis.

8. What Next?

There is fairly universal agreement on several aspects of the Internet:

- It is, and will remain, essentially global in its reach.
- Many new mechanisms for access and delivery are coming into use. XDSL over telephone wires, cable TV, DBS satellite and mobile wireless services are just a few.
- No carrier will have a complete monopoly; interactions between different carriers in the same technology, and between different technologies, will be essential.
- At the moment there is a vast amount of surplus transmission capacity coming on-stream as carriers tool up to introduce the new fibres. While traffic may eventually grow to fill them up, this will take some time.
- There is still a huge instability in the organisational structure of the carriers. There will be further international mergers both geographic and between different types of carrier; there will also be further break ups of carriers who now have monopolies both in their own local area and between different technologies.

In view of the above, it should be possible to extend the national testbeds which are currently being set up into a significantly larger scale. Not only should single carrier testbeds be provided, but also multiple-carrier ones. The US experience has shown that once the research community is let loose on such testbeds, they enrich it, encourage it to become more rugged, and provide unforeseen new uses. The resulting increase both of traffic and functionality is of value to all the participating organisations - even if the carriers individually are competing.

The heterogeneity of the participating network technologies should be greatly enriched. Many technologies should be included such as DBS, mobile wireless, cable and xDSL. While one view of these last is that they are only access technologies, their integration into the global system are non-trivial. How one does do resource location, routing and access control into a French uplink from a US host has not an obvious answer. Considerable work needs to be done to understand how to integrate roaming using mobile telephone access in different countries with multicast conferencing. The requirements for gateways with proper Quality of Service for Video-on-Demand with different types of xDSL are not clear. Any of us who remember the problems in making simple ISDN terminals work in different countries - for a supposedly internationally standardised service - must understand the problems to be expected with the future Internet services, which are much more complex. What will happen when there is a large growth in the use of IPv6 - with IPSEC - over the global Internet? How will Certification Authorities, set up according to the rules of the different countries, be usable as part of distributed, multi-national applications?

There will always be problems in the globalisation and the scaling of many of these services. However, many of the problems will only surface when testbeds are set up on a large scale with an interesting mix of such services. It is not enough to carry out such testbeds on a national scale - even if the national scale is as large as Japan or the US. Many problems can be resolved on a national scale; many others will only occur when some of the assumptions held nationally will be found to be at variance with those in other countries.

One problem that the European Union has had in setting up testbeds under the Framework programmes was the desire to set them up to cover every country in the European Union. The differing attitudes of the various carriers made this an expensive exercise - and led to a minimal set of services being provided. The charges demanded often did not really admit that the services were often experimental, the research project customers would not have paid for the services commercially, and far less than commercial grades of availability were required (or provided). Moreover only a minimal set of fixed terrestrial services were actually provided. From the carrier viewpoint, there has always been concern that any rates established might become a precedent. To the extent that subsidised services were offered, there was also a concern that these would infringe some of the emerging Competition regulations of the national and European Commission regulators. A quite different approach is required both for the sake of the research projects which are to be supported, and to develop the high quality, and varied, infrastructure that are both technically possible and desired by the customers.

9. A Specific Proposal

9.1. The Proposal Itself

I propose that from the suppliers', governmental, research network providers' and users' viewpoints, a new initiative be launched. The regulators should agree that the provision of subsidised facilities for the provision of advanced Internet services be permitted. With the present technology, it is comparatively simple to derive Virtual Networks from the *networks to support research*. Many of those countries not yet encouraging the provision of *research networks,* should derive such VNs from their *networks to support research,* or should provide them in some other way. There should be active encouragement for research projects, whether involving commercial or academic researchers, to use these networks for advanced services - not only for their development but also user trial up to some level. There should be encouragement to carriers to connect novel services to these networks - including mobile, satellite and cable.

National carriers should become willing to connect their services to these networks. They should admit that it is not only the perfection of their specific services that matter, but also their integration into the larger scheme of things. Provided it is for this class of service, special rates should be permitted. Ideally the whole cost of participation should be regarded as a Research and Development activity, thereby benefiting from the tax credit or other fiscal means provided locally to support this class of activity - with some protection from international World Trade concerns.

Many countries are already doing the things mentioned in the previous paragraphs; many are not. We should, however, go further. Carriers should also be encouraged (and be prepared) to provide international connectivity to other similar national networks. Where carriers are multinational, they should be encouraged to link into several national networks. Where several carriers operate in a particular country, they should all be encouraged (and persuaded) to participate.

Clearly the actual interconnection of some of the services mentioned would be both difficult and experimental. There should be active encouragement of suppliers in working with these projects to provide experimental facilities - again not on a fully commercial basis.

It is important that the providers of the facilities and services get some return for their effort. Part of that return will be non-financial but in terms of experience. Even this depends on a proper level of usage for advanced purposes. Many of the countries which already have programmes to encourage advanced network research or utilisation should announce projects in those areas that the suppliers also feel there is need for action. Where there are already projects spanning several countries, there should be specific measures to encourage international working. At present this means, too often, the countries working with US projects like Internet-2; sometimes, as with the whole mobile or DBS fields, there would be much to be gained from more regional collaboration.

The users must also sometimes relax their prejudices. For example a research network provided on a non-commercial basis may have poorer availability than a commercial one. Idealistically it is desirable for all encryption to be strong encryption; for the sake of these research applications, some compromise may be needed here for some applications or countries to encourage a wider provision of the requisite infrastructure.

For this sort of initiative to work, many ingredients are need. It is necessary for the funders of national programme to consider the international dimension; the NSF has made some important initiatives in this regard in some of its latest programmes. The EU Framework programmes have the right words in their Programme Plan, and even the right elements in the Call for Proposal due to be announced in September 1999. They include the phrases testbeds, integration, international, Internet, Mobile, Satellite. Whether there is the scope to oil the wheels of the sort of programme I am outlining here is another question. Most funding bodies like to fund particular, well defined, specific projects. While the Programme clearly will be made up of such projects, they may be too difficult to define *a priori* to get easy funding. Only too often there will be uncertainties at the time a proposal deadline is required; only if the funding bodies buy into both the detailed proposals and the spirit of the programme will such projects get funded.

9.2. Specific Milestones

Bob Kahn got it right when he put forward the slogan "A gigabit to the desktop" at the start of the Gigabit Testbed project. While I actually disagreed that this was the right goal, and it probably was not even intended, it was the right message to appeal to those who needed to support the project - both in industry and government. I will set a more modest set of milestones.

- At least 15 countries, including ones in Asia, Australia, the Americas and Europe agree to participate by the end of 2000.
- At least 10 terrestrial carriers agree to participate by the end of 2000, and six companies to provide routers, switches, multiplexors, management systems.
- By the end of 2000, at least three mobile carriers, three each of DBS operators, LEO satellite providers and TV cable service operators agree to participate.
- At least OC-12 available on the major trans-oceanic routes by the end of 2001 - with at least OC-3 available optionally to some national nets.
- By the end of 2001, a testbed service allowing 20 person, high band, multicast, secured VoD and multimedia conferencing using mobile terminals at a few hundred Kbps in four countries, terrestrial terminals at 10-20 Mbps over normal LANs, and both xDSL over the telephone and cable modems into the VPN Internet. The system will use IPv6, QoS, IPSEC, Mobile IP and have a rugged multicast infrastructure. It will be watched by a further 400 in ten countries (including some with relatively poor communications infrastructures) with limited audio feedback through the normal Internet.

10. Conclusions

We have reviewed some of the experiences with testbeds in the past, and considered how they are developing. We have developed that a wider definition of testbeds is needed, and that they should both be international in extent and more far-reaching in technology. We have outlined some of the steps that need to be taken to provide open, global, available and heterogeneous Internet facilities.

We have discussed the ideas outlined here with some carriers, national network operators and equipment suppliers. Several have expressed strong interest and are keen to help the proposal get started. Let us hope that all of you and your sponsors will buy into the concept.

REFERENCES

[can] http://www.canarie.ca/index.html

[dfn] http://www.dfn.de/win/

[int2]http://www.internet2.edu/presentations/Overview-DVH

[kir99] Kirstein, PT: Early Experiences with the Arpanet and Internet in the United Kingdom, IEEE ann. Hist. Comp. 21, 1, 38-44, 1999.

[ngi]http://www.darpa.mil/ito/research/ngi/index.html

[quan] http://www.dante.net/quantum/

[surf] http://www.surfnet.nl/surfnet-publicaties/