

LONG READ: HOW TO DESIGN A FIBRE OPTIC NETWORK

It would be nice to be able to call in Harry Potter to conjure up your new fibre-optic network. But alas, instead it requires some incredibly detailed design work. In addition to the cost, there are many parameters to take into account, namely how much redundancy you think you can afford, how fast the network should be, where the nodes will be, what equipment to choose, the type of fibre and how the network will be physically built to actually achieve the redundancy.

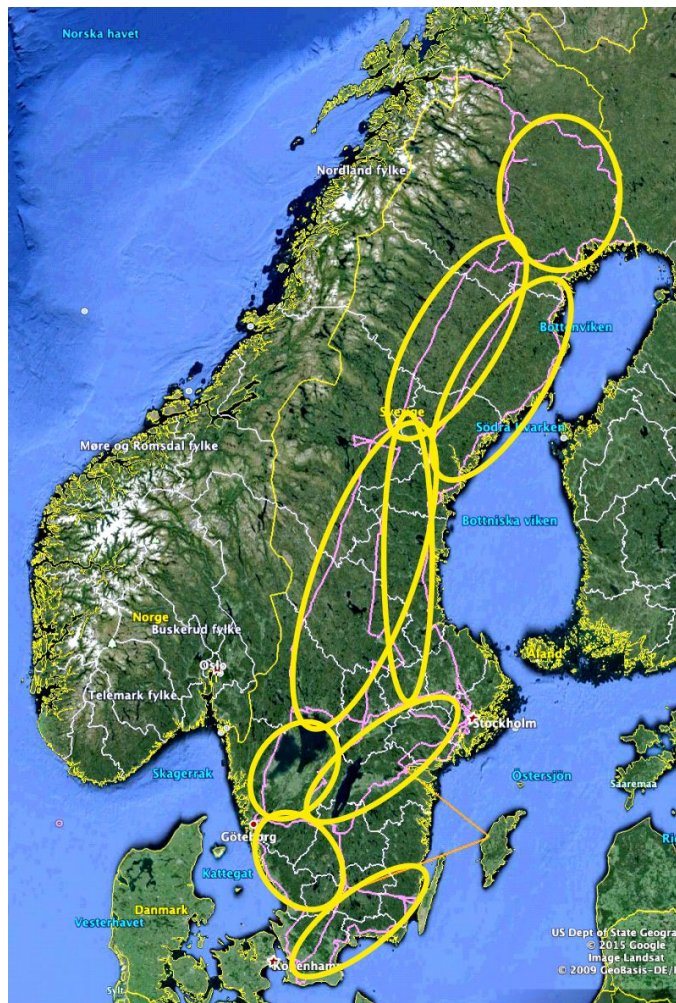
For this, the operative manager Börje Josefsson and his co-workers had to frown their foreheads and do some seriously extensive work. But how did they do it?



Börje Josefsson, presenting a wrinkle-free forehead.

– You begin by collecting a group of people with serious knowledge. Everyone isn't required to know everything, though. You also try to find external people with great knowledge, to ask for good advice. The SUNET workgroup included said Mr. Josefsson, Per Nihlén and Magnus Bergroth. A good example of an external person was e.g. Peter Löthberg. All these have long-standing experience in network design, and were all involved with the design and procurement of our present network, the OptoSunet, Börje starts.

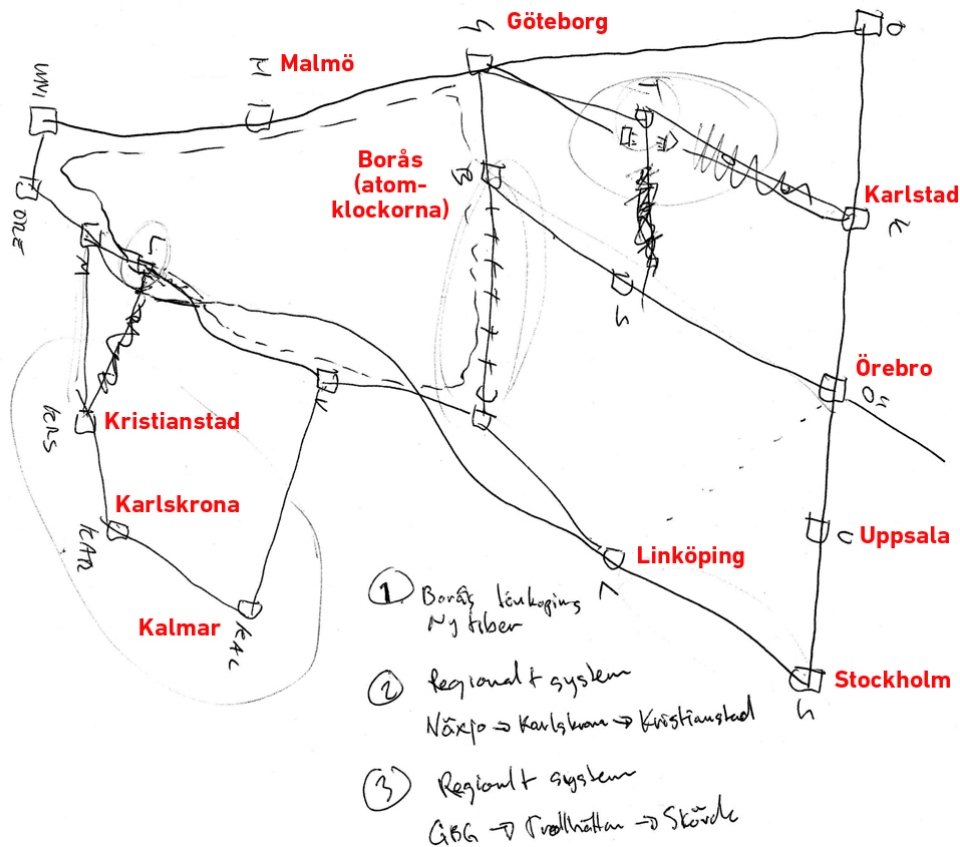
But having all this good old knowledge is also a dilemma. Technology isn't standing still. OptoSunet looks the way it does because of the technological limits that prevailed around the year 2006. As we designed the OptoSunet, about ten years ago, there was a 1000 kilometre limit to transmitting an optical signal before it had to be regenerated (tidied up) using electronics, which is expensive. These limits are now gone. We might span up to 2500 or even 3000 kilometres before regenerating. There are more degrees of freedom to create another layout than ever before.



The main rings of the new ring structure

A long time ago we knew that we wanted more local routing, and more redundancy, if possible. We have tried to accommodate this with the new topology. The OptoSunet topology was a star in which all data had to be routed through Stockholm. You might see the new network as a series of rings. Or you might view it as a snowman with a number of snowballs on top of each other. The ring structure makes us more resilient to diggers cutting the fibre, and similar catastrophes. Most regions will have three pathways, or at least one northbound and one southbound.

Thanks to our decision to avoid a star topology and the fact that you might span 2000 kilometres without regeneration, plus the fact that there is lots of competition on fibres these days, our new fibre network is cheaper than the previous one. We need fewer kilometres of fibre but will have three paths, where we previously had only two. We expect to be running ten times as fast, with better redundancy, but the network won't cost more.

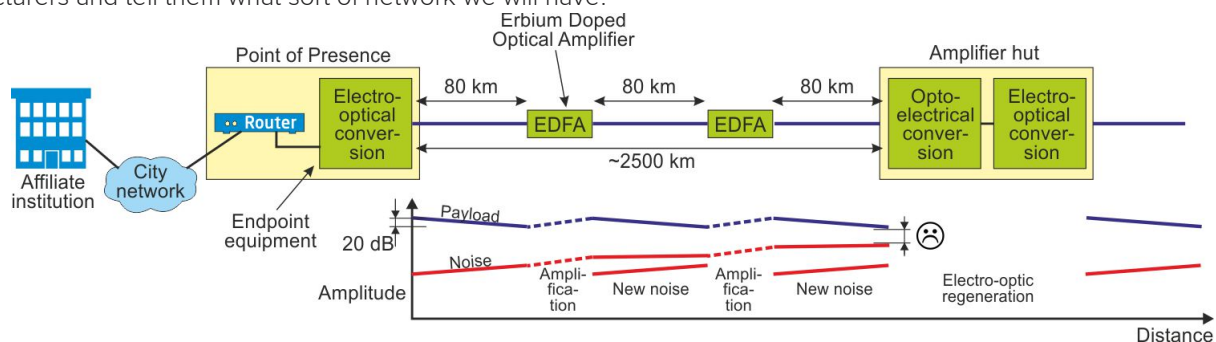


All great projects start as a sketch on a napkin or on the back of an envelope. The sketch above is the beginning of the design of the southern part of the network, drawn by hand on a piece of paper someone snatched from a printer. The author has made some efforts of decoding it. For example, note the text “Borås, atomklockorna” pointing to the city of Borås where the main atomic clocks keeping track of Swedish Mean Time resides.

KEEPING YOUR STUFF IN ORDER

Ideally, when designing a network one should have all the facts on the table. But you don't. Instead, you very soon will find yourself in a Catch-22 situation. The first thing an equipment supplier will ask is what sort of fibre you have. Before April 2015 we had no fibre. To present the fibre supplier with some reasonable demands, you need to know the demands of the equipment supplier.

This evil circle must be broken. We decided to write some approximate fibre specifications and then procure a fibre network, to actually get to know which fibre we had. After the procurement finished, in April 2015 we now know which fibre paths we have, the distance between the in-line amplifiers, the dispersion, and so on. Now we are able to talk to the equipment manufacturers and tell them what sort of network we will have.



Because the fibre is not perfectly transparent, the optical signal will be attenuated en route. Modern fibre technology tolerates some 20 dB attenuation before the signal needs amplification. It will happen sometime after 80 kilometres. At 80 kilometre intervals, all-optical EDFA:s (optically pumped erbium amplifiers) are inserted. On longer spans, the EDFA can be supplanted by a more advanced amplifier type, a so called Raman amplifier, in which the fibre itself acts as the amplifying


component. (But think of it. 80 kilometres is quite fantastic. How thick can you make an ordinary window pane before it gets opaque? Half a metre? 3 centimetres is enough to get 3 dB's attenuation, that is, half the light is lost. The same happens after 14 kilometres in an optical fibre!)

The problem with primitive optical amplification is that noise will creep up on the way. Finally, so much noise has added up, that the signal needs to be interpreted, changed into electrical bits and bytes and then re-converted back into light. That's regeneration. As regeneration is expensive, one wants to avoid it if at all possible. Fortunately, Sweden isn't 2000 kilometres long, so using present technology, no electro-optical regenerating amplifiers will be needed.

One goal with our design was to create a network with as little noise as possible, to get the maximum OSNR (Optical Signal to Noise Ratio). Nothing very significant seems to have been happening in optical amplification in the last ten years, and we hope this will not change in the future. If we use the best stuff available, we should be able to live with it for the whole of the network's calculated lifetime.

Everyone agrees that we will have to change endpoint equipment during the network's lifetime though, but just avoiding changing optical amplifiers is quite something. That's why we spend quite some time to work out the ideal combination of Raman and EDFA, just to get the lowest possible noise figure. And it's no good having too many amplifiers either, as they add noise, too.

POINT OF PRESENCE



Fiberslangar i gatan/Fiber conduits in the street.

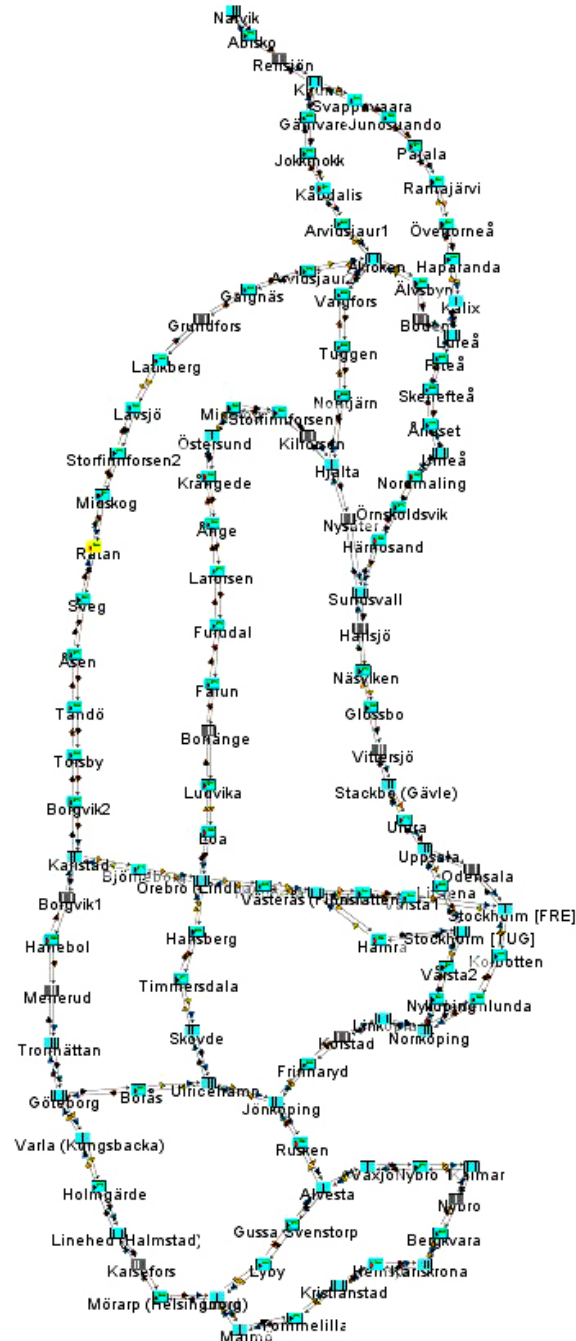
Bild/image: Bidgee, CC BY-SA 2.5

Data traffic is normally switched in and out of the network in a university city. The exchange should never happen on a college, as one SUNET affiliate must never be dependent on any other affiliate, such as during e.g. a power cut. Instead, the exchange takes place in a POP.

The fibres must be utilised to the maximum. This means they will be lit up with light of different wavelengths and will simultaneously carry several data streams of 100 Gbps. The combined traffic in and out of a POP will probably be in the range of a terabit per second.

The traffic from the POP to the university is usually carried through a city network, where SUNET is unable to determine the fibre's exact location. Instead, Tele2 decides this. The only requirement is that the connection be redundant. SUNET will still be the organisation lighting up the fibre.

REDUNDANCY



Den logiska kartan över SUNET C. Ringarna är klart urskiljbara. Notera att Gotland saknas./The logic map of SUNET C. The rings are obvious. Please note that Gotland Island is missing.


The “red” and “green” optical fibres of today’s OptoSunet run together between e.g. Stockholm and Uppsala, although hopefully on different sides of the road. But they are still in the same part of the country. If both fibres between Stockholm and Uppsala are cut, redundancy is lost.

With the new design, Uppsala will have one fibre pathway going north and another going south. Somewhere around Sundsvall the fibre turns westward into the country and then back down to Örebro. And if that pathway is cut, data may

continue to travel further north and turn back from Luleå and get back via Karlstad. This will make the path some 15 milliseconds longer, but we would rather have that than complete breakdown.

Simultaneously, today's absolute Stockholm-centric dependence will be eliminated. If the traffic is routed in Luleå or Malmö, the probability of a total blackout is much smaller than with today's topology.

Also, note that the network ends in Narvik in Norway. The idea is to bring Norwegian university traffic down to southern Norway through Sweden. The same thing goes for the Swedish power grid. Norway has a whole lot of fjords, making it difficult to put power pylons along the coastline. Northernmost Norway is nothing but "coastline". You could probably lay down a fibre on the sea bed along the coast, but it could easily be torn by ship anchors, when storms start blowing from the Atlantic. Instead we choose to hang the fibre in the power grid.



Optisk fiber av självbärande typ, observera stålvejern i mitten.
/Self-supporting optical fibre. Note the steel wire in the centre.
Bild/Image: Srieffler, CC BY-SA 3.0

To continue with the power grid, one may ask why Tele2 and SUNET thinks it's such a good idea to hang the fibre in the Swedish power grid. The fibre could easily be dug down in the ground, but then a farmer may go there and dig it up and break it. Hung in the power grid, it is out of everyone's reach. Someone might fire at it with a shotgun, but it really doesn't happen that often. Also, it is very rare for a 400 kilovolt pylon to fall over in a storm. So far it has never happened.

On the other hand it might be a problem during cold weather when all power lines are needed to transfer power to the whole country. In such case Svenska Kraftnät will not shut down a line, to allow Tele2 to send a technician up a pylon to repair something. It is possible to use self-supporting fibre, which is not part of the overhead ground wire, which may be repaired while the power is on. The same type of self-supporting fibre could be used in the regional networks, to be hung alongside 10 kV and 20 kV lines. It stays intact even if a tree should fall on it.

A more serious problem is that the fibre is sensitive to wind. As it shakes back and forth, the angle of polarisation of the light inside the fibre will spin. The data transmission is dependent on the polarisation, and if it spins too fast, the connection may be lost. We would like to return to this subject in a future article.

AT THE UNIVERSITY

A few universities believe they have enough capacity already. They will not be hit by any extra costs. The equipment SUNET will install on their premises will be plug compatible with their present equipment. Should they, however, wish to increase their capacity to 100 Gbps, they will have to upgrade.

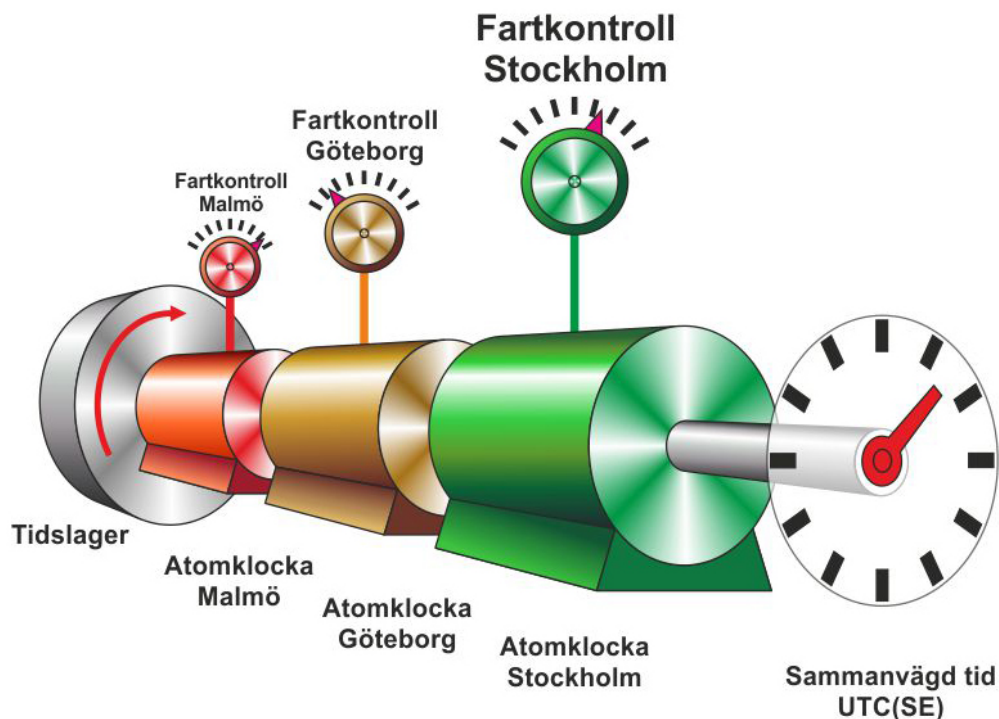
A number of smaller colleges still uses 1 Gbps. They are asked to upgrade to 10 Gbps at an early stage. It is simply cheaper to connect them in this way, than to try to get some sort of exotic hardware that will be able to manage the age-old 1 Gbps. The way our topology looks now, it will not be cost effective to downshift SUNET's 100 Gbps one hundred times.

MULTIPROTOCOL LABEL SWITCHING AND OTHER SPECIAL SERVICES

MPLS is a way to create a logical (virtual) network on top of a physical one. This is presently used on SUNET. The virtual network functions as a separate network and may be used for private point-to-point links.

More and more universities start co-operations and merge with each other. One example is the Uppsala University that merged with the College of Gotland Island more than a year ago. The college became an internal institution in the University and because of this the University wanted the network at Campus Gotland to be on equal footing with their other institutions, with access to the intranet etc. That's why SUNET has a virtual link between Uppsala and Gotland which makes Gotland look like it is behind the Uppsala firewalls. Other colleges in Sweden utilise the same service.

We do more or less the same for researchers who need a special connection out of Sweden, for some special research task. Some examples of this may be the Onsala radio telescope at Råö having their own wavelength in our fibres, and SP Technical Research Institute in Borås (those guys who fetch Swedish Mean Time to their atomic clocks from the Bureau of Weights and Measures in Paris).



SUNET will also take part in a service of great benefit to society as a whole. When the NTP Project is finished, SUNET will act as a link for Swedish Mean Time between a number of atomic clocks at various places in the country. The illustration above shows how a number of atomic clocks, likened to motors with their own speed controls, will cooperate with the SP Technical Research Institute in Borås, likened to a large flywheel, to maintain Swedish Mean Time (sammanvägd tid UTC(SE)) with a high level of accuracy. This is needed if Sweden should lose connection with the BIPM in Paris for an extended time. The robustness and protection against denial of service attacks will be greatly improved, compared to the present situation.

As always, SUNET will henceforth attempt to satisfy any researcher's need for connections between any two places in the country.

QUALITY OF SERVICE

– We do have Quality, and we have Service, but we don't know what Quality of Service is, Börje says jokingly. We simply have enough bandwidth to not having to quality grade traffic. If so, who would get priority? No one is able to say whether KTH or Chalmers are more important than any other affiliate.


A curiosity: A long time ago, during the discussion about procuring firewalls for the KTH, someone put the clever question: Which way should the firewall work? Are we supposed to protect KTH from the world, or protect the world from KTH? We still have more or less the same problems.

NUMBER OF AFFILIATES

A whole lot of organisations are connected to SUNET today. There are a total of 34 universities and colleges, and some of them are big-time bandwidth users. Let's just mention Chalmers, whose radio observatory, the Råö Telescope is part of several radio astronomical projects, requiring data streams in the order of 50 Gbps. Or one might mention another set of big-time users: the student housing networks. These networks are home to some very big data gobblers, consuming limitless amounts of gigabytes.

There are also art institutions of various types, such as the Museum of Architecture, The Army Museum, The Museum of Ethnography and The Air Force Museum. Another one is the open air museum Skansen, which will probably not up its bandwidth significantly within a foreseeable future.

Then there are 33 “other organisations” among which we find the real data cannons, which will require a lot of network uptime, real soon. One of them is the new neutron gun, the European Spallation Source (ESS) in Lund, which will need a massive link for data transmission to Europe. On the other hand, institutions like CERN in Switzerland will start sending massive amounts of data to Swedish universities and colleges, when the new search for dark matter starts in the LHC accelerator.



Mörk materia (blå) i universum. Kanske. Om den finns. Det återstår att ta reda på./Dark matter (blue) in the universe. Maybe. If it exists. We'll find out, for sure.

Sweden has a total of six supercomputing centres, commonly referred to as SNIC (Swedish National Infrastructure for Computing). Some of them are located in Linköping, Stockholm and Umeå. They carry out computation for a variety of projects, such as particle research at CERN, meteorology for SMHI and more. This makes them top consumers of bandwidth.

The SMHI Meteorological Institute is connected to its colleagues at MET in Norway, and they use SUNET, NORDUnet and Norwegian university Uninett to exchange meteorological data between their nodes.

WHAT'S HAPPENING?

– I'm not sleepless over the time schedule so far. Right now (May, 2015) we need to hurry and decide what endpoint equipment to use, finish that part of the design and make sure we stay within the economic limits, Börje continues. We have to place an order for hardware some time during the autumn, as the delivery time will be some three to four months.

FUTURE AND LOGISTICS

As the new SUNET C is airborne in the second half of 2016, everyone will get 100 Gbps. But 200 and 400 Gbps are beckoning at the horizon. Looking 15 years ahead, terabit speeds may be the norm. This will necessitate new endpoint equipment, whereas the in-line amplifiers may be retained. This is in itself a great cost saving. For the same reason, anyone will be able to change to 200 Gbps transponders where needed, without disrupting any other network activity.

Now the fibre contract has been signed, and Tele2 is driving the network expansion. The great logistical challenge will appear when the endpoint equipment arrives. We will have to install hardware at about 100 sites. SUNET will require help from many different organisations. Tele2 will have to help, because they own the sites. The hardware supplier will have to deliver to the proper places, and electricians must be allowed in to connect the power. Finally, technicians from SUNET need to get there to configure the hardware.

New users will be connected all the time. The EISCAT 3D is an ionospheric research project located in the north of Scandinavia, which is just starting up. The idea is to observe and image the ionosphere (aurorae, solar storms etc) in three dimensions, to increase understanding of space weather in general. The various EISCAT sites need network connections. In general, this means getting 10 Gbps to Karesuando, Porjus and Abisko. That is, fibre in the alpine world.

When the new SUNET C is finally up and running, all the equipment from the old OptoSunet must be removed, and may be sold off or disposed of in other ways. This will be a challenge, too, not least logistically.

What will the traffic patterns look like in the future? No one knows. We can try to look at what the patterns are today, but then suddenly someone will start a new bandwidth-sucking project which puts everything on its head. For this, we have no forecasts, Börje Josefsson concludes.

FURTHER READING

Only in Swedish. Please use Google Translate.

This is how the Swedish Internet works: <http://techworld.idg.se/2.2524/1.399148/omojligt-att-stoppa-sa-fungerar-sveriges-internet>

How to build a future-safe city network: <http://techworld.idg.se/2.2524/1.586008/sa-byggs-ett-stadsnat-for-framtiden>

The fibre that does 3.4 terabits per second: <http://techworld.idg.se/2.2524/1.599956/34-terabit-i-sekunden>

The topology of the global Internet: <http://www.sweclockers.com/artikel/19339-sa-ar-det-globala-internet-uppbyggd>

Skriven av



JÖRGEN STÄDJE

Jag heter Jörgen Städje och har skrivit om teknik
och vetenskap sedan 1984. Friskt kopplat, hälften
brunnet!