Framework for path finding in multi-layer transport networks

Dijkstra, F.

Publication date
2009

Document Version
Final published version

Abstract

In only a few decades the use of computer networks has dramatically increased. Today, networks are ubiquitous in society: they are used for surfing, email and financial transactions. The capacity of the network has increased so much, that it is now possible to transfer massive data sets that recently were shipped on tape or disk by a courier. Vast data streams in nuclear physics, radio astronomy and recent transmission of movies in exceptional high quality are only a few of the prime examples of these transfers.

The Internet is not always suitable to transmit large amounts of data between a limited number of locations. In some cases it is better to create a dedicated network connection for a specific application. A dedicated connection is better if it is cheaper or if the data streams are so large that they would disrupt the regular Internet traffic. Besides the already mentioned examples of transfer of scientific data and movies, the data exchange between two sites of the same company can be one of those applications. In most other cases the regular Internet is the better choice.

Since 2005 national research networks such as CAnet in Canada and SURFnet in the Netherlands have started offering lightpaths to their customers, besides their regular Internet connectivity service. These are dedicated connections between only two locations, as opposed to the regular Internet, where every host can connect to all places in the world. The name lightpath comes from the optical fibre networks that often provide these connections. A network that provides both lightpaths as well as regular Internet connectivity over the same physical infrastructure it is called a *hybrid network*.

Telecom providers have offered dedicated network connections since a long time, although the term *lightpath* is relatively new. While lightpaths mimic leased lines in technology and capacity, their dynamics resembles a telephone connection. Ideally, lightpaths can be provisioned automatically, without intervention from a human network operator.

Lightpaths can span multiple administrative domains. A typical network connection between two universities firsts crosses a campus network, then a national research network, then goes through an international peering to another research network, then another campus network and finally to an internal network within a building. Different persons and organisations administrate all these networks.

Each of these networks has been built and deployed at different times, and each network designer will have made different choices from the list of available technologies. Each network consists of different layers, and one network may be able to switch at the Ethernet layer, another at the SONET layer and a third may be able to switch different colours of light over a fibre.

It turns out to be very relevant that lightpaths cross multiple networks that are capable of switching at a different network layer. Different choices in technology for each network lead to potential incompatibilities. This thesis shows that these incompatibilities make the problem of finding a shortest path through the networks significantly more complex than the problem of finding paths through networks where the switching occurs at the same layer, such as on the Internet. It turns out that paths in hybrid networks have restrictions that depend on choices made elsewhere in the path. The restrictions of a path through a network with a single switching layer are independent of the choices elsewhere in the path.

This thesis shows that paths that cross multiple technologies, the so-called multi-layer network connections, can go in loops. It is possible that a shortest path traverses the same fibre twice. In addition, it turns out that a segment of a shortest path does not have to be a shortest path in itself.

The largest part of this thesis covers the exploration and description of a formal model for multi-layer networks, including hybrid networks. The presented model is technology independent. This is an essential feature, since it means that neither the model nor the path finding algorithms need to be adjusted as new technologies are invented in the future.

Besides the model, the results of this work are two comparable multi-layer path finding algorithms. One of these algorithms has been implemented, and this thesis closes with a description of its first use.