Downloaded from UvA-DARE, the institutional repository of the University of Amsterdam (UvA)
http://dare.uva.nl/document/164781

File ID	164781
Filename	Chapter 7: Summary and conclusion

SOURCE (OR PART OF THE FOLLOWING SOURCE):

Type	Dissertation
Title	A semantic model for complex computer networks : the network description language
Author	J.J. van der Ham
Faculty	Faculty of Science
Year	2010
Pages	viii, 154
ISBN	978-90-814160-2-3

FULL BIBLIOGRAPHIC DETAILS:
http://dare.uva.nl/record/331711

Copyright

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other then for strictly personal, individual use.

UvA-DARE is a service provided by the library of the University of Amsterdam (http://dare.uva.nl)

Chapter 7

Summary and Conclusion

The provisioning of an inter-domain lightpath is currently still a complicated process. The user must formulate a request to the provider. The provider then checks availability in the network, and must figure out a path through the inter-domain network. Each of the domains then check their availability in turn, and if the destination is reachable, the lightpath is provisioned. Once the lightpath is working, it is delivered to the end user.

The different parties involved in the provisioning process mostly communicate ad hoc through telephone and email. This allows many opportunities for mistakes, or misunderstanding, complicating the provisioning and the monitoring processes. This thesis provides contributions towards facilitating the inter-domain network provisioning process. We have divided this in two parts, first creating a clear and inter-operable description of the network, and second to examine the impact of aggregating topology information on the performance of pathfinding.

In the first part of this thesis we have examined the problem of describing computer networks. In chapter 2 we have described requirements for creating interoperable descriptions of computer networks. We have used these requirements to evaluate existing ways of representing computer networks. None of these representations satisfy all of the requirements.

In chapter 3 we have presented our Network Description Language (NDL). The beginning of the chapter describes the first version of NDL, which provides a simple ontology to describe network topologies. While the vocabulary is small,

it is already very powerful.

Based on our experiences with the first version of NDL and based on related work on GMPLS, we have extended NDL to describe multi-layer networks. After extensive study of the ITU-T G.805 standard we have concluded that it is possible to create a technology-independent description of multi-layer networks with NDL. Inspired by the theory in ITU-T G.805 we have created multi-layer NDL. Multi-layer NDL provides a vocabulary to independently describe technologies. These technology descriptions can then be referenced from multi-layer topology descriptions.

In chapter 4 we have shown several of our applications that show new possibilities with the use of NDL. We have gained valuable experience and feedback from our development of applications and tools to support and manage real-world networks using NDL. Our choice for RDF was first put to the test when we developed generators and validators. Descriptions in RDF can be very verbose, but our generators and validators proved to be simple and adequate solutions. Furthermore, our Python NDL Toolkit (pynt) has lowered the boundary of developing applications using NDL. The toolkit has also allowed us to rapidly prototype ideas and has been a valuable input for the continued improvement of NDL.

In summary, NDL provides a distributed information model for the description of topologies for inter-domain pathfinding. So this allows us to give a positive answer to the first research question: *It is possible to create a distributed information model for the description of topologies for inter-domain pathfinding.*

Network operators do not always wish to share their full topology, either for security, business, or for scalability reasons. On the other hand it is necessary to share some degree of topology data to enable inter-domain lightpath planning. A solution is to publish aggregated versions of the network topologies. In the second part of this thesis we have examined the impact of aggregation on the performance of inter-domain pathfinding.

In chapter 5 we have described different methods for aggregating domain topologies, the Full Mesh, Star and Single Node strategies. We have summarized three earlier publications on pathfinding in aggregated topologies. Unfortunately these studies use different measures for performance, and are therefore hard to compare. Even when comparing the general trends, they show some conflicting results on the performance of the Star strategy.

In order to better quantify the performance impact of topology aggregation in optical networks, we have performed our own experiments. Chapter 6 describes our experimental set up, and shows the results of putting different strategies to the test. This has allowed us to accurately determine the impact

of different aggregation strategies on inter-domain pathfinding.

Based on our results we can give an affirmative answer to the second research question: *topology aggregation does indeed have an impact on pathfinding, albeit not a great impact.*

If the Full Mesh strategy is used, and the topology description is often updated to reflect the current availability, then there is almost no difference in performance compared to having all the available information. Of course, a Full Mesh aggregation that is often updated maintains all the information relevant for inter-domain pathfinding. If no updates are performed, then the Star aggregation should be used. The pathfinding will suffer an impact from the aggregation, but the Star strategy performs better than a Full Mesh without updates. The Single Node strategy performs worst, even worse than the Full Mesh without updates. The experiments and our analysis of the results show the way that topology aggregation has an impact on inter-domain pathfinding, providing an answer to the second research question.

7.1 The Road Ahead

7.1.1 RDF Infrastructure Descriptions

In our view the next step in network descriptions is the definition of ontologies that can cover the whole end-to-end infrastructure: from the actual content being distributed to the Storage elements holding the data to the CPUs rendering the images and the display to visualize it. Figure 7.1 illustrates the concept of RDF representation of all this elements forming the overall infrastructure.

Our vision is that a media content locator will be able to consume the descriptions of all the architectural components that form the end-to-end infrastructure. From this information it will build the optimal paths from the storage elements to the visualization displays, making the inter-domain lightpath provisioning a piece in the overall orchestrated effort. The use of RDF allows easy extension of the topology and domain schema with other schemata.

Our vision is that an application will be able to consume the descriptions of all the architectural components that form an end-to-end infrastructure. This information includes computing resources, storage resources, visualization resources, network resources, content descriptions, et cetera. All resources can be linked with loose couplings to allow a meta-scheduler application to orchestrate all resources together in a combined effort [7].

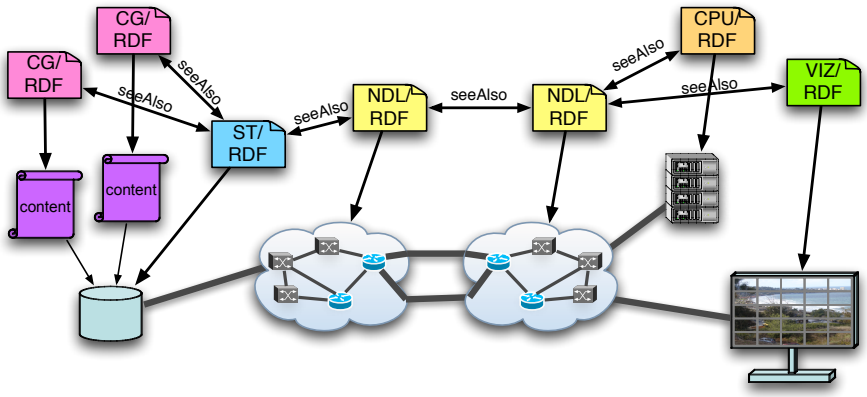

Figure 7.1: *RDF Infrastructure Descriptions*

7.1.2 Topology Aggregation

In our experiments we have not used a crank-back algorithm, however we expect it to give similar results as the Full Mesh view with updates. The crank-back algorithm provides a way of using information about false-positive failures to do recalculation. Effectively this allows a requester to get the same level of information for path computation as with constant updates. The difference is that the main burden of computation is then moved from the domain to the client. Another point for future research is inserting a delay in the propagation of domain updates, so that we can confirm earlier results of Awerbuch et al.

Another open issue is mixed aggregation strategies, our experiments used the same aggregation strategy for all domains. In practice different domains will make different choices for topology aggregation. It could be investigated what effect this will have on the performance of the inter-domain network pathfinding.

Multi-layer pathfinding in optical networks is currently a topic of research[70]. For now, this pathfinding is based on full topology information, since this is a very hard problem in itself. An open issue is defining aggregation strategies for multi-layer topologies. It is far from trivial to apply aggregation to multi-layer topologies. Besides the connectivity information that is aggregated, the encoding and the adaptation capabilities must also be considered.