Link Loss Tomography and Topology Synthesis

Rhys Bowden

Thesis submitted for the degree of

Doctor of Philosophy

in

Applied Mathematics

at

The University of Adelaide

Discipline of Applied Mathematics

School of Mathematical Sciences

January 22, 2015

Contents

Abstract

Accurate and timely performance data are of vital importance for network administration. However, modern networks are so large and transmit such enormous quantities of data that collecting the complete measurements can be wildly impractical. *Network tomography* uses the measurements that *are* available to infer underlying performance statistics. In this thesis we consider estimating average packet loss rates on links from more easily collected path measurements. Most such work has concentrated on tree-networks, but here we consider the problem on a general network where the problem is typically underconstrained. In that context we need some criteria to select a solution from the infinite set of possibilities. Here we exploit the compressive sensing assumption of sparsity. However, although the assumption of sparsity makes a great deal of sense in this context, the standard conditions required for results in compressive sensing theorems do not hold for realistic routing matrices. We show that despite this, the underlying techniques can still provide useful answers. What's more, we show that the apparently inconvenient structure of routing matrices can actually help in solving the problem efficiently. We provide CTD, a new algorithm for finding sparsest solutions that is orders of magnitude faster than one of the standard compressive sensing algorithms, and which provides more certainty. We also provide a version of CTD for working with a limited number of measurements, CTDn.

The success of a tomography algorithm often depends upon the underlying network topology. To test CTD we use two sources of topologies: the Internet Topology Zoo and synthetically generated topologies. We propose a new method for topology synthesis, Combined Optimisation and Layered Design (COLD), that mirrors the real-life design process of a data network. Since real data networks are designed, they are in some sense optimized to fulfil a function. However, strict mathematical optimisation is rarely performed at the router lever; rather the PoP-level is typically the one being optimized, because it is both more stable and less intricate. Once the PoP-level network is determined, the router-level network can be created using a templated design process, increasing regularity and aiding management and debugging. COLD mirrors this process by having two layers: PoP-level optimisation of realistic economic constraints subject to demand; and then router level synthesis with a templated design process

using graph products. We show COLD produces sensible, varied yet controllable output with a clear relationship to the relatively few input parameters.

We test CTD and CTDn with both Topology Zoo topologies and COLD-generated topologies. Because the COLD process can be controlled to generate a wide variety of topologies this allowed us to test the sensitivity of CTD and CTDn to the form of the underlying topology. We show that CTD and CTDn perform well across a wide variety of topological structures and forms of measurement, while relying on less detailed information than previous approaches.

Declaration

I certify that this work contains no material which has been accepted for the award of any other degree or diploma in my name in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text. In addition, I certify that no part of this work will, in the future, be used in a submission in my name, for any other degree or diploma in any university or other tertiary institution without the prior approval of the University of Adelaide and where applicable, and partner institution responsible for the joint award of this degree.

I give consent to this copy of my thesis, when deposited in the University Library, being made available for loan and photocopying, subject to the provisions of the Copyright Act 1968.

I also give permission for the digital version of my thesis to be made available on the web, via the University's digital research repository, the Library Search and also through web search engines, unless permission has been granted by the University to restrict access for a period of time.

Signature.............................. *Date*............

Acknowledgments

Many thanks to my supervisors, Professor Matthew Roughan and Professor Nigel Bean, for their time, for their excellent advice and expertise, for their patience and forbearance, and for being examples of the kind of academic I can only hope to be one day.

Thanks to my colleagues and collaborators. I've enjoyed working with all of you, and I hope we work together again in future.

Thanks to Mum and Dad for their support and love throughout my whole life, not just during my Ph.D. candidature.

Thanks to all my family and friends for making life so enjoyable.

And finally thanks to Jess for everything. Including being Jess.