Position Statement
Workshop on Internet Routing Evolution and Design (WIRED) 2003

Aman Shaikh
Computer Engineering Department
University of California
Santa Cruz, CA 95064
aman@soe.ucsc.edu

1 Routing at the Box-level

Surprisingly, routing research has focused very little on the routers themselves. I believe the following issues need more research attention:

- **Router design:** The research community is not aware of a whole spectrum of issues related to the design and internals of a router. Some examples are: what hardware and software issues vendors have to grapple with, how the protocols are implemented, how protocols contend with one another for resources, and how forwarding tables are constructed/updated. Addressing these issues require vendors to open up by de-mystifying and providing more insight into the black-box world of routers. Making router internals accessible to the broader research community without giving out proprietary information in itself is a challenging task.
- **Black-box characterization of routers:** This would help in developing better analytical and simulation models, and vendor-independent benchmarking methodologies.

2 Routing at the Network-level

Understanding the routing system at the network-level mostly boils down to collecting and analyzing routing data from the networks. Specifically, following issues need to be addressed:

- **"Smart" collection of data:** It is not feasible to collect data from every router in the network. For link-state protocols like OSPF collecting routing updates from a few vantage points provides a network-wide view of OSPF routing state and is probably sufficient for most purposes. On the other hand, for path-vector protocols like BGP, collecting routing updates from any vantage-point generally gives only local view of BGP routing state from that vantage point. Does this mean that we have to collect BGP updates from all the routers? Generalizing this, how can we gain maximum visibility and understanding of routing system by collecting the "right" kind of data from the "right" set of vantage-points?
- **Making data accessible to research community:** As networks are instrumented to collect routing data, we need ways of making this data accessible to the research community without compromising the proprietary nature of the data. This may prove to be a significant technical challenge.
- **Joining data from multiple sources:** Understanding interaction of routing protocols often requires joining of multiple data sources. Do we need new methods and tools for joining these data sources?

3 Routing System as a Whole

Most of the routing research has so far focused on individual protocols in isolation. Moving forward, the community needs to focus on the routing system as a whole. Consider this example. Improving IGP convergence has received a fair bit of research attention lately. However, IGP convergence is only a part of the overall convergence process for (tier-1) ISPs since most of the traffic in such networks is actually routed by BGP (which in turn depends on the IGP for egress-point selection). Routers from one of the vendors have a BGP scan process that periodically updates BGP

next-hops in response to IGP routing changes. Since this period is 60 seconds by default, the IGP convergence time is a small portion of the overall convergence time in a lot of instances. The bottom-line is that one needs to focus on both IGP and BGP to reduce the overall convergence time.

As a part of looking at the routing system as a whole, we also need to look at how different layers (IGP, IBGP/MPLS, EBGP) of routing and the interface between them needs to evolve. I believe that understanding of how protocols interact at box and network levels will form the basis for understanding the routing system as a whole.

4 Analytical and Simulation Models

There is a dearth of good simulation and analytical models of the routing system today. Having sound models and simulators based on these models is crucial for the following reasons:

- To analyze large, network-wide dynamics that are difficult to study in a lab setting or in a real network. For example, performing scalability analysis of routing protocols would be hard without good simulators.
- To analyze trade-offs between various design alternatives — both at the box and network-level.
- To understand the (potential) impact of new features and extensions as they are added to the routing protocols.

The box and the network-level understanding of the routing system will hopefully lead us to better analytical and simulation models than what we have today. We also need methodologies and tools to validate the models.

5 Reducing the Pain of Administrators

The research community has mostly ignored network administrators' pain by assuming that the network core is simple and all the complexity resides in the end-systems. This is not true. Network cores are incredibly complex, and a significant part of this can be attributed to routing. Routing protocols were not designed with manageability in mind. There is a lot of scope for research in the management of routing systems. Some of the issues that need to be addressed are:

- **What's in a network?** Often administrators do not even know basic questions like how many routers they have in the network, how they are physically connected, what routing protocols they run etc. Often such knowledge exists as bits and pieces in the minds of a few individuals. As researchers, what abstractions, methods and tools can we develop to help administrators understand their networks better?
- **Configuration management:** At present, network configuration is done at a device-level: every router is configured separately. We need to move from such device-level configuration to network-level configuration. This requires developing languages for expressing network-level design goals, and tools for translating them into device-level configuration commands. These days routers tend to have a plethora of configurable parameters; often operators resort to intuition as well as trial-and-error for setting these parameters. How can we turn the "art" of setting configurable parameters into a "science"?
- **Debugging and trouble-shooting routing problems:** Finding the root-cause when something goes wrong in a network is often a pain-staking task. Although routing protocols are designed and (hopefully) implemented with isolation in mind, in practice they often tend to interact with one another in extremely weird ways. As a part of data collection and analysis exercise mentioned earlier, we need to devise ways of correlating various message-streams in *real-time*, to help operators identify root-causes of the problems quickly and reliably. Another challenge that we need to address is trouble-shooting problems that span multiple domains of control.
- **Hardware and software upgrades:** Seamlessly upgrading the routing hardware and software is not an easy job. Several network meltdowns have occurred as undesirable side-effects of such upgrades. Is there something we can do to make upgrades pain-free? Can we devise ways of modifying/extending the existing protocols or even completely replacing these protocols without too much bloodshed?

Acknowledgments

I am thankful to Jennifer Rexford, Zhuoqing Morley Mao and Renata Teixeira for their feedback and comments which went a long way in improving this statement.