Towards a Modular Routing Control Platform for Inter-domain Routing

Yi Wang
Princeton University
yiwang@cs.princeton.edu

Jennifer Rexford
Princeton University
jrex@cs.princeton.edu

Reality 1: It's hard to make changes to the routing architecture, especially to inter-domain routing. As a result, although everyone knows that BGP is not perfect, we are still running it (with more and more patches/extensions), and may keep on running it for quite long.

Reality 2: No matter how hard we try, we are always building tomorrow's legacy systems.

Noting these two realities, we feel that it's important to build some routing system that can provide both of the following functions:

- Encourage innovations to routing (given Reality 2)
- Ease the (architectural or protocol) transition from present to future (given Reality 1)

Control Plane Needs Programmability

We think the control plane, where routing decisions are made, is the right place to combine both of the functions. However, in today's IP networks, control plane functions are distributed at individual routers. The lack of a network-wide view at the control plane not only introduces huge complexity and vulnerability in configuration, but also makes it difficult to achieve some network level goals (e.g., traffic engineering), not to mention the support for above two high level functions.

The intertwining of control plane and data plane at the router level is a pain and largely unnecessary. As a result, there has been a series of work refactoring of the AS-level routing control, especially about the boundary between control/management plane and the data plane elements [1, 2, 3]. The RCP (routing control platform) is designed to control inter-domain routing in a logically centralized way so that it can select routes on behalf of routers with a complete view of the network. It is also backward compatible [1].

This is all well and good, but still not enough. By the nature of routing control, any individual AS should be able to define its own network-level objectives and functionalities. To best support this goal, we argue that the RCP should not only provide configurability, but also provide *programmability*, which provides the maximum support for extensibility.

Modularity Provides Programmability

Not surprisingly, people have thought about using modular design to make routing systems more extensible. The Click modular router is designed to be assembled from packet processing modules called elements, which make it highly flexible and configurable [4]. XORP provides inherent extensibility by a composable framework of routing processes, each of which in turn is composed of modular processing stages [5]. Most recently, the CONMan project proposes a new design of a network management architecture that reduces the management complexity (analogous to 4D's discovery and dissemination plane) by introducing module abstractions and defining simple generic interfaces between data-plane modules and control-plane modules [6].

Click and XORP provide extensibility to the data and control planes respectively, but both in the context of a single router. The same extensibility is still missing in the inter-domain routing control/the RCP context. What we propose here is to make the RCP modular and programmable so that people can:

- Program existing modules (e.g., the decision making logic)
- Implement new function modules and "plug and play"
- Compose different modules together to achieve some specific functionality desired by a particular AS

Getting the Interfaces Right is the Key

Achieving such modularity requires careful design to make a modular routing control platform (MRCP) truly

extensible. A key question to answer is "what are the right interfaces?" To answer this question, we first need to define the functions a routing control platform needs to provide. Today, inter-domain routing is expected to do more than just selecting AS paths. We expect the MRCP to also provide support to expressive business relationship configuration, traffic engineering, and various BGP extensions.

How to combine all these functions in a modular and extensible way? An MRCP is essentially an event (message) driven system – once a routing message arrives at the MRCP, it will trigger a sequence of processing, where each step is at a function module (policy enforcement, decision making, etc). We believe that every function of an MRCP should be a module for two major reasons: first, this offers the most extensibility – nothing is irreplaceable; second, this makes the reimplementation or replacement of a specific function easy. This doesn't necessarily mean any module is optional. For example, any MRCP must have one decision making module and (at least) one policy module, although such module can be replaced by another with similar function. It is the responsibility of the implementer of a module to make it comply the defined interfaces. It is the responsibility of the core of an MRCP to (dynamically or statically) link/plumb different modules together.

One noteworthy thing is that each module may need information other than the routing message itself, which may be from policy configuration, or traffic engineering configuration, or side information such as traffic monitoring results, etc. Hence we essentially need two types of interfaces:

- Internal interfaces between MRCP modules, which are mostly used by programmers
- External interfaces, which are exposed to other systems/programs or network operators.

Getting these two types of interfaces right would make an MRCP very extensible.

In summary, we argue that the control plane needs programmability. We believe that modularity is a good way to provide programmability, and with careful interface design the MRCP (modular routing control platform) is a promising way to both encourage routing innovations and to bridge present and future inter-domain routing architectures/protocols/algorithms.

References

[1] N. Feamster, H. Balakrishnan, J. Rexford, A. Shaikh, and K. van der Merwe, "The case for separating routing from routers," in *Proc. ACM SIGCOMM Workshop on Future Direction in Network Architecture*, August 2004.

[2] A. Greenberg, G. Hjalmtysson, D. A. Maltz, A. Myers, J. Rexford, G. Xie, H. Yan, J. Zhan, and H. Zhang, "A clean slate 4D approach to network control and management," *Computer Communications Review*, vol. 35, pp. 41–54, October 2005.

[3] J. van der Merwe, A. Cepleanu, K. DSouza, B. Freeman, A. Greenberg, D. Knight, R. McMillan, D. Moloney, J. Mulligan, H. Nguyen, M. Nguyen, A. Ramarajan, S. Saad, M. Satterlee, T. Spencer, D. Toll, and S. Zelingher, "Dynamic connectivity management with an intelligent route service control point," in *Proc. ACM SIGCOMM Workshop on Internet Network Management*, September 2006.

[4] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek, "The Click modular router," *ACM Transactions on Computer Systems*, vol. 18, pp. 263–297, August 2000.

[5] M. Handley, O. Hudson, and E. Kohler, "Xorp: An open platform for network research," in *Proc. ACM SIGCOMM Workshop on Hot Topics in Networking*, October 2002.

[6] H. Ballani and P. Francis, "Conman - taking the complexity out of network management," in *Proc. ACM SIGCOMM Workshop on Internet Network Management*, September 2006.