Central Control Over Distributed Routing – Public Review

Teemu Koponen
Styra, Inc.
koponen@styra.com

Modern network management is foremost a case of precise control: to implement practical traffic engineering and network security policies, a network operator needs to control the exact path traffic takes through a network. While conceptually simple, accomplishing this level of control is surprisingly difficult with traditional network control planes; in the end, the underlying distributed routing protocols provide the operator only with indirect controls over the traffic forwarded. Indeed, it was this inadequate control that motivated SDN and its wholesale departure from these proven, but distributed protocols.

This paper explores a less disruptive path to SDN. Contrary to the canonical SDN design leveraging OpenFlow (to manage switch state) and distributed systems techniques (to synchronize controller state), the authors build on existing and familiar routing protocols to synchronize state, critically allowing centralized management of *unmodified* IP routers. The authors observe that given desired forwarding state per a router, it is feasible to derive a fake link state topology that– once injected to the link state protocol by a controller– causes each unmodified routing stack to locally compute and install the desired forwarding state. The resulting design is technically sound. It is deployable, scalable, and failure tolerant but the backwards compatibility comes with an inherent trade-off: its expressivity is still confined to expressivity of IP routing. A controller cannot program forwarding state matching beyond the destination address.

It would be trivial to follow the paper with a debate about its long-term practical value, *i.e.*, whether a standard routing protocol (perhaps further improved with minor extensions) could, or should, serve as an SDN switch management protocol. The debate would undermine the true value of this paper. The paper reminds us that SDN is not architecturally about a particular wire protocol but about decoupling the control and data plane, and how in fact there are multiple ways, with varying trade-offs, to achieve it. Hopefully, this paper will foster more out-of-the-box research on the unexplored trade-offs in the fundamentals of SDN.