

Merlin: Programming the Big Switch

Robert Soulé¹, Shrutarshi Basu², Robert Kleinberg², Emin Gün Sirer², and Nate Foster²

¹University of Lugano, souler@usi.ch

²Cornell University, {basus,rdk,egs,jnfoster}@cs.cornell.edu

Introduction. Software-defined networking (SDN) makes it possible to treat an entire network as a single switch that forwards traffic between its ports [5]. This “big switch” abstraction presents programmers with a global view that hides the complexities inherent in the physical network: distributed state, complicated forwarding rules, and device-specific configuration. But although it is an appealing abstraction, an key question remains: *how should we effectively program a big switch?*

Existing SDN programming languages [2, 7, 9, 1, 8] suffer from limitations that make them unable to adequately capture the big switch abstraction—they either focus on forwarding or force programmers to express policies in terms of hop-by-hop processing steps. Furthermore, these languages do not support policies that involve applying richer functions (e.g., implemented on middleboxes, end hosts, or custom hardware) to certain packets.

This paper presents Merlin, a new network programming language designed to address three essential aspects of big switch programming: (i) dividing traffic across multiple sub-policies; (ii) controlling forwarding paths; and (iii) provisioning bandwidth in terms of limits and minimum guarantees. Merlin provides intuitive constructs for specifying each of these features, as well as a compiler that maps source programs into constraint problems whose solutions determine allocations of resources such as paths and bandwidth. The Merlin run-time allows allocations to be dynamically adjusted, and provides mechanisms for verifying that updated allocations obey the constraints specified in the original program.

Merlin Policy Language. As an example to illustrate, consider the program shown in Figure 1(a). It places a bandwidth cap on FTP control and data transfer traffic, while providing a bandwidth guarantee to HTTP traffic. Syntactically, the program consists of a sequence of statements, followed by a logical formula. Each statement contains a predicate on packet header fields that identifies a set of packets, a regular expressions that describes a set of forwarding paths, and a variable that tracks the amount of bandwidth used by packets processed with the statement. The statement on the first line, with variable x , asserts that FTP traffic from host `192.168.1.1` to `192.168.1.3` must travel a path that includes deep-

packet inspection (`dpi`). The next two statements identify and constrain FTP control and HTTP traffic between the same hosts, respectively. Note that the FTP control statement does not include a `dpi` constraint in its forwarding path, while the HTTP statement includes both a `nat` and a `dpi` constraint. The formula on the last line declares a bandwidth cap (`max`) on the FTP traffic, and a bandwidth guarantee (`min`) for the HTTP traffic. Overall, the Merlin language facilitates direct expression of high-level policies, without worrying about how those policies will be enforced in the underlying physical network.

Network Provisioning. The Merlin compiler implements three tasks: (i) it translates the global program into one or more locally-enforceable policies; (ii) it determines forwarding paths, places virtual network functions, and makes bandwidth allocations by mapping policies to constraint problems; and (iii) it generates the low-level instructions needed to realize the program using the switches, middleboxes, and end hosts.

First, the compiler rewrites the formula so that each constraint applies to packets at a single location. Given a formula of one term with n identifiers, the compiler produces a new formula of n terms that collectively imply the original. By default, the compiler divides bandwidth equally among the traffic classes, although other schemes are permissible. The `max` term in the example is localized as `max(x, 25MB/s)` and `max(y, 25MB/s)`. Next, the compiler encodes policies into constraint problems that can be solved to determine forwarding paths through the network, placement of packet-processing functions on individual devices, and allocations of bandwidth. Figure 1(b) illustrates the fragment of the encoding for the HTTP statement above. The red path in the rightmost graph indicates a path that satisfies all constraints. Merlin supports several path-selection heuristics, including one that optimizes for shortest paths and others that minimize the allocation on any link. Although our prototype uses the Gurobi optimizer [3] to find a solution, Merlin is not dependent on any particular algorithm. For example, the abstractions provided by the language could be mapped to an approximation algorithm instead. Finally, after the compiler has determined a placement for the system components, it generates the appropriate code to enforce the

