Programmable Data Planes

COS 597E: Software Defined Networking

Jennifer Rexford
Princeton University
MW 11:00am-12:20pm

Data Plane

- Streaming algorithms that act on packets
  - Matching on some bits, taking a simple action
  - … at behest of control and management plane
- Wide range of functionality
  - Forwarding and access control
  - Buffering, marking, shaping, and scheduling
  - Rewriting header fields (e.g., NAT)
  - Traffic monitoring and deep packet inspection
  - Encryption, compression, and transcoding

A Need for Speed

- High link speed
  - 10 Gbps, 40 Gbps, 100 Gbps
- Small packets
  - 40-byte TCP ACK packets
- Small time per packet
  - 40 Gbps =124 Mpps with 320-bit packets
  - 8 ns to process a packet
- Routers need low latency
  - Relatively limited opportunity to batch or pipeline

A Range of Technologies

- ASIC (App-Specific Integrated Circuit)
  - Fast, dense chip, but expensive to change
- FPGAs and network processors
  - Fast, reconfigurable, hard to program
- Graphics Processing Units
  - Massive parallel computation on small cores
- Software on commodity computer
  - Easy to program, but I/O bandwidth, memory copying, and interrupt overheads

Click Modular Router

Click Motivation

- Flexibility
  - Add new features and enable experimentation
- Openness
  - Allow users/researchers to build and extend
  - (In contrast to most commercial routers)
- Modularity
  - Simplify the composition of existing features
  - Simplify the addition of new features
- Speed/efficiency
  - Operation (optionally) in the operating system
  - Without user needing to grapple with OS internals

Router as a Graph of Elements

- Large number of small elements
  - Each performing a simple packet function
  - E.g., IP look-up, TTL decrement, buffering
- Connected together in a graph
  - Elements inputs/outputs snapped together
  - Beyond elements in series to a graph
  - E.g., packet duplication or classification
- Packet flow as main organizational primitive
  - Consistent with data-plane operations on a router
  - (Larger elements needed for, say, control planes)

7

Push vs. Pull

- Packet hand-off between elements
  - Directly inspired by properties of routers
  - Annotations on packets to carry temporary state
- Push processing
  - Initiated by the source end
  - E.g., when an unsolicited packet arrives (e.g., from a device)
- Pull processing
  - Initiated by the destination end
  - E.g., to control timing of packet processing (e.g., based on a timer or packet scheduler)

8

Click Language

- Declarations
  - Create elements
- Connections
  - Connect elements

```
src :: FromDevice(eth0);
ctr :: Counter;
sink :: Discard;

src -> ctr;
ctr -> sink;
```

- Compound elements
  - Combine multiple smaller elements, and treat as single, new element to use as a primitive class
- Language extensions through element classes
  - Configuration strings for individual elements
  - Rather than syntactic extensions to the language

9

Handlers and Control Socket

- Access points for user interaction
  - Appear like files in a file system
  - Can have both read and write handlers
- Examples
  - Installing/removing forwarding-table entries
  - Reporting measurement statistics
  - Changing a maximum queue length
- Control socket
  - Allows other programs to call read/write handlers
  - Command sent as single line of text to the server

10

An Observation…

- Click is widely used
  - And the paper on Click is widely cited
- Click elements are created by others
  - Enabling an ecosystem of innovation
- Take-away lesson
  - Creating useful systems that others can use and extend has big impact in the research community
  - And brings tremendous professional value
  - Compensating amply for the time and energy ☺

11