

OVERVIEW

- EDS control networks require dependable delay guarantees on delivery of network packets
 - Packets must be delivered between hosts with *guaranteed upper bounds on end-to-end delays*
 - Should be resilient to flooding or other network level denial-of-service (DoS) attempts
- Traditional approaches (e.g., CAN [1], AFDX [2]):
 - Not suitable for commercial off-the-shelf (COTS) systems
 - Require expensive custom hardware and software
 - Proprietary, complex, expensive
- Our approach:
 - Leverages the capabilities of **software-defined networking (SDN)**
 - Reduces the management and integration overheads
 - Guarantees timing constraints for traffic in hard real-time systems
 - Can easily be integrated with COTS SDN hardware
 - Isolates flows into different queues
 - Provide *stable* quality of experience (e.g., end-to-end delays) even in the presence of heterogeneous (e.g., best-effort) traffic

SOFTWARE DEFINED NETWORKS

- Communication network components:
 - Control planes* (decide packet forwarding rules)
 - Data planes* (perform the actual actions on packets)
- SDN simplifies access to the network configuration
 - Controller*:
 - Logically centralizes the control-plane state
 - Switch*:
 - Contains a table processing pipeline and a collection of physical ports

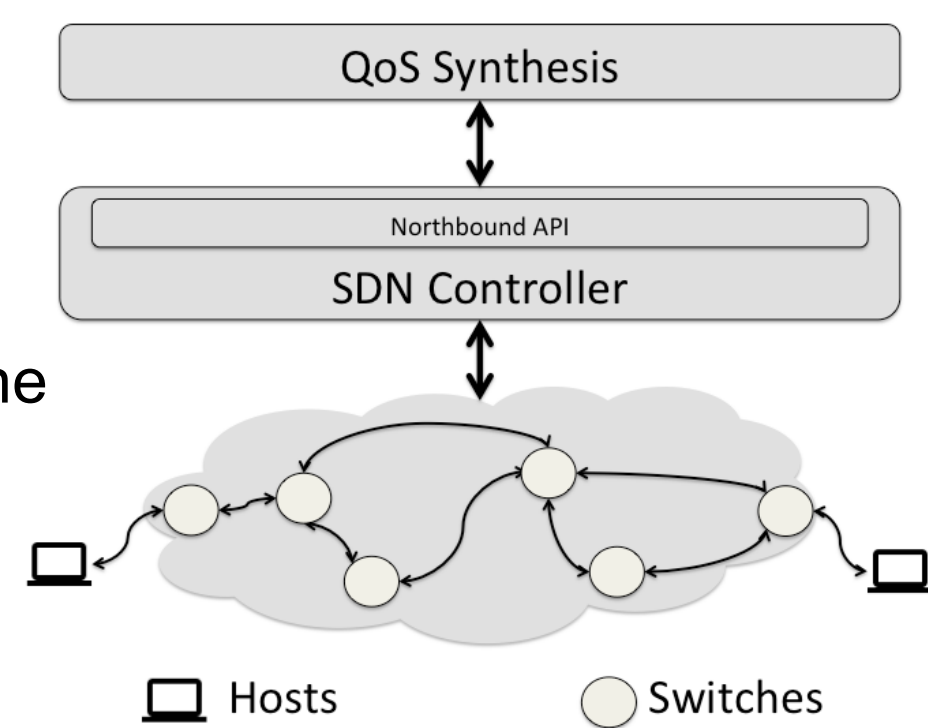


Figure: An SDN with a six switch topology

PROBLEM OVERVIEW & SOLUTION SKETCH

- EDS control networks often include *high priority/critical* traffic flows
 - Essential for the correct and safe operation of the system
 - Have stringent timing requirements
 - Can tolerate *little to no* loss of packets
 - Examples*: sensors for closed loop control and actual control commands in power grid systems

How to develop mechanisms to *guarantee end-to-end delays for high-criticality flows* on networks constructed using SDN switches?

- Given an SDN topology:
 - Switches and controller
 - A set of real-time flows with specified delay and bandwidth guarantee requirements
 - The problem is to
 - Find paths for the flows and
 - Map the flows to the queues of switch ports
 such that the *end-to-end delays can be guaranteed* for the maximum number of critical flows
- Our proposal:
 - Develop a *criticality-aware constrained* path selection algorithm
 - Allocate *each flow to an individual queue*
 - Intuition:
 - Overprovision the bandwidth
 - Critical real-time flows do *not* experience queueing delays even in the presence of increased packet flow on non critical flows

PATH LAYOUT AND FLOW MAPPING

- The selection of an optimal path for each flow subject to delay and bandwidth constraints is NP-complete
 - Extended existing heuristic [3] and developed *polynomial* algorithm
- Isolating flows using *separate queues* results in lower and more stable delays
 - Especially when traffic rate in the flow approaches the configured maximum rates




Figure: The measured mean and 99th percentile per-packet delay for the packet the two-switch, four-host topology

EXPERIENCE AND EVALUATION

- Experimental setup:
 - Mininet [4] topology with Open vSwitch [5] configured switches
 - Ryu [6] SDN controller
 - netperf [7] to generate the UDP traffic
 - Network with Heterogeneous flows
 - [1, 5] critical flows along with [1, 3] non-critical flows
- Results and Observations:
 - Our flow rules and queue configurations isolate the critical flows from the non-critical traffic
 - Non-critical flows do *not* affect critical flows
 - The mean and 99th percentile delay experienced by the real-time flows *always* meet their delay requirements.




Figure: (a) The empirical CDF of (a) average, (b) 99th percentile delay with different number of flows. For each set of flow $f \in \{2, 5\}$, total $f \times 25 \times 5$ packet flows (each for 10 seconds) were examined

For our experiments, we do not observe any instance for which a set of schedulable flow misses its deadline (i.e., packets arriving after the passing of their end-to-end delay requirements)

CONCLUSION AND FUTURE WORK

- Presented mechanisms
 - that provide end-to-end delays for critical traffic in EDS networks using COTS SDN switches
- Hence, future EDS control networks can be better managed, less complex (fewer network components to deal with) and more cost effective
- This initial effort can be extended in several directions:
 - Multiplex flows and yet meet their timing requirements
 - Impose *admission control* policy
 - Prototype and evaluate on actual *hardware switches*

REFERENCES

- M Farsi, K Ratcliff, M Barbosa. An overview of controller area network. Computing & Control Engineering Journal, 1999.
- ARINC Specification 664, Part 7, Aircraft Data Network, Avionics Full Duplex Switched Ethernet (AFDX) Network. 2003.
- S. Chen and K. Nahrstedt. On finding multi-constrained paths. In *IEEE ICC*, 1998.
- B. Lantz, B. Heller, and N. McKeown. A network in a laptop: rapid prototyping for software-defined networks. In *ACM SIGCOMM Workshop on Hot Topics in Networks*, 2010.
- Open virtual switch. <http://openvswitch.org/>.
- Ryu controller. <http://osrg.github.io/ryu/>.
- Netperf. <http://www.netperf.org/netperf/>.