Slicing 5G fronthaul networks using programmable switches

Nishant Budhdev, Raj Joshi, Pravein Govindan Kannan, Mun Choon Chan, Tulika Mitra

National University of Singapore

ACM CoNEXT '20, Barcelona, Spain

Slicing in Cellular Networks

Slicing allows operators to slice a physical network into multiple virtual networks

Each virtual network can be used for different

- Use cases such as eMBB, mMTC, uRLLC
- Mobile Virtual Network Operators (MVNO)

Existing works on slicing focus on different parts of the cellular network such as

- Wireless Spectrum: Orion
- Radio Access Network (RAN): SoftRAN, FlexRAN, PRAN
- Core Network: Softcell, MobileFlow, KLEIN

There exists no work on slicing the fronthaul!

Fronthaul Slicing

Fronthaul connects the frontend base station to the RAN and carries digitized radio signals between the two parts

Fronthaul slicing enables a distributed RAN architecture which has the following benefits

- improved RAN average and worst-case performance
- reduced CAPEX and OPEX

Enabling fronthaul slicing requires multipoint-to-multipoint networks which support

- point-to-point (one enhanced Radio Equipment Control (eREC) — one Radio Equipment(RE))
- point-to-multipoint (one eREC — several REs)
- multipoint-to-multipoint (several eRECs — several REs)

System Design

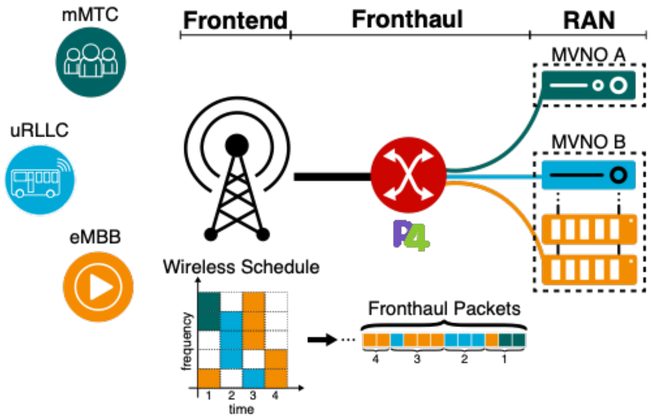

"All packets in the fronthaul can be predicted *in advance* using the wireless schedule."

Each wireless schedule has the following properties:

- generated every millisecond
- contains at most 10 users

To store the forwarding information we use SRAM-based memory (called register arrays) in the dataplane of programmable switches

Ring buffer is created using register arrays with 2 indexes

- **write index**: used by "Schedule Packet" to add entries in the ring buffer
- **read index**: used by Fronthaul Packets to read entries in the ring buffer

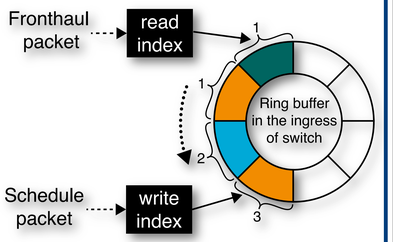

Evaluation

Design implemented on Barefoot Tofino switch using ~1000 lines of P4 code

Overall latency for adding 10 entries in the SRAM is 3.5us

- adding 10 entries in the switch requires less than 2µs
- latency for generating and transmitting the packet to the switch is 1.5µs

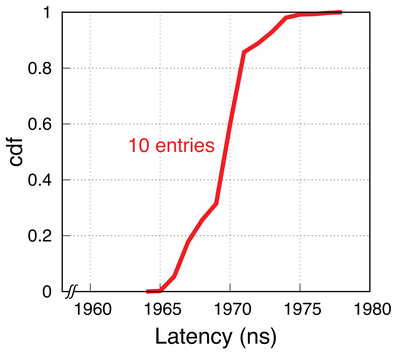

Future Work

Extend system design to support

- packet reordering and drops
- packet prioritization

Perform comprehensive system evaluation to test

- real world performance using network schedules captured from commercial operators
- support for anyhaul traffic

Support new time synchronization paradigm independent of master(eRE) and slave(eREC) design

Acknowledgements

This work was supported by the Singapore Ministry of Education Academic Research Fund Tier 1 (Grant Number: T1 251RES1910).