interdigital

XRC: An Explicit Rate Control for Future Cellular Networks

Moteza Kheirkhah
InterDigital

Mohamed Kassem
University of Surrey

Gorry Fairhurst
University of Aberdeen

Mahesh Marina
University of Edinburgh

IEEE ICC 2022

Motivation

- Most of today's cellular end-to-end data flows are originated from TCP
- TCP performs poorly over cellular links due to rapid fluctuation.
- The extent of the fluctuation may even be worse with emerging high capacity wireless technologies (mmWave)
- As emerging to 5G, TCP would even be worse:
  - emerging high capacity wireless technologies (mmWave)
  - applications and services are expected to vary significantly, from very low-latency to very high-bandwidth
- TCP is not designed to meet the requirements of these diverse applications, especially when they are mixed at a shared cellular link

Motivation (cont.)

- Several end-to-end cellular-specific congestion control schemes have been proposed (e.g., Verus, Proprate, ExLL, Raven, Copa, PCC Vivace)

- These schemes outperform TCP but are sub-optimal in being unable to accurately and quickly track capacity changes (even BBR failed to do it)

- Tackling such problems is fundamentally difficult with a purely end-to-end congestion control mechanism

Performance of TCP variants over wireless networks

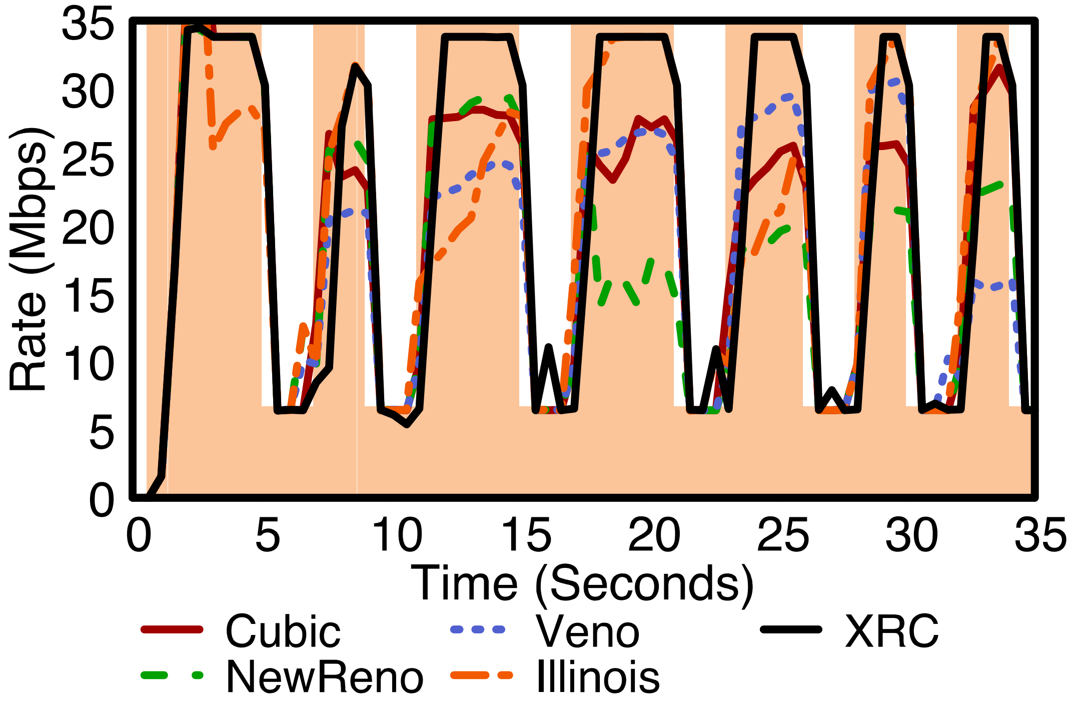

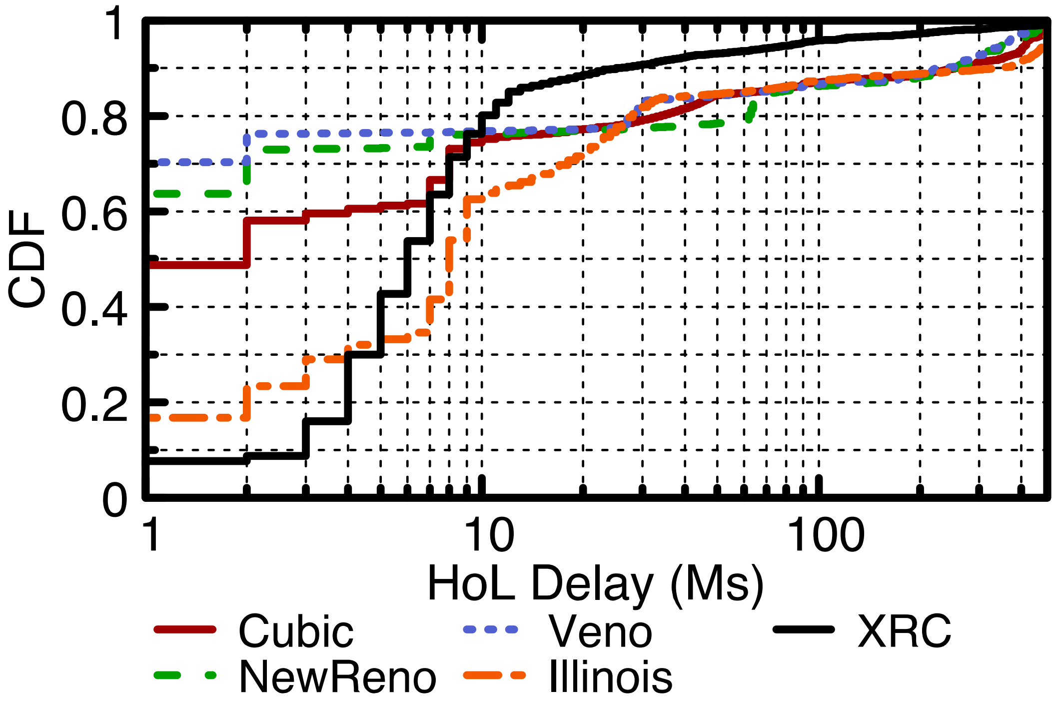

Common TCP variants fail in both utilizing the radio resources as well as controlling the queuing delay

eXplicit Rate Control (XRC)

- An adaptive TCP congestion control that exploits explicit feedback from the cellular's base station.
- XRC requires two pieces of information from the base station:
  - (1) the rate available to UEs' active flows
  - (2) the Head-of-Line (HoL) delay of the base station's RLC buffer
- XRC can also distinguish the Internet bottleneck from the cellular bottleneck and react accordingly
  - By constantly monitoring the RLC's HoL delay and the end-to-end queuing delay

XRC goals

- Full utilization of scare radio resources
- Low queuing delay (i.e., essential to meet the low-latency traffic requirements)
- Fair co-existence between heterogeneous competing flows

XRC system components

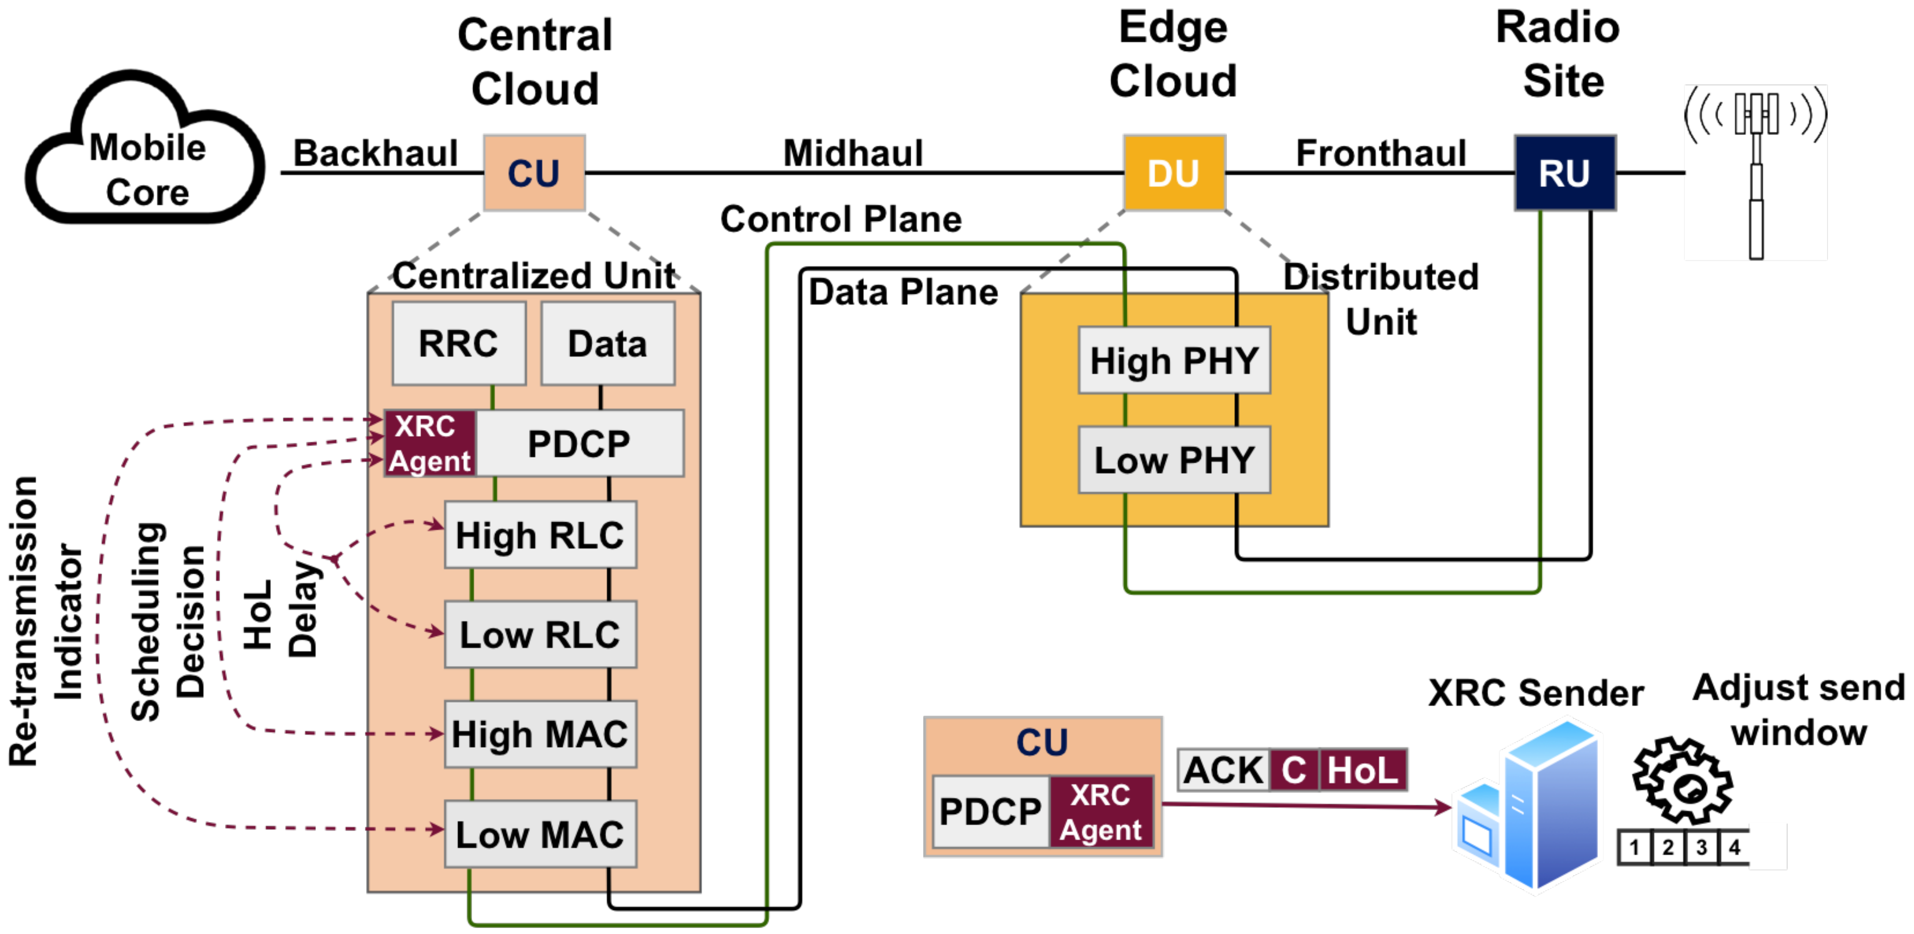

XRC congestion control algorithm (1)

- Upon reception of an ACK with explicit feedback, XRC calculates a target congestion window (twnd) as follows:

$$twnd = R_{exp} * RTT_{min} \qquad (2)$$

- The XRC Sender estimates the queuing delay (***σ***) in every window of data as the maximum of the average queuing delay and the RLC HoL delay (**HoL_{exp}**) as follows:

$$\sigma = \max\{\sum_{j=1}^{N} RTT_j * \frac{1}{N} - RTT_{min}, HoL_{exp}\} \qquad (3)$$

XRC congestion control algorithm (2)

- **Congestion Avoidance**
  - When cwnd **<** twnd and queueing delay (σ) **<** queueing budget
    - Increase cwnd rapidly until it reaches twnd within 2 RTTs
  - When cwnd **>=** twnd and queueing delay (σ) **<** queueing budget
    - Increase cwnd gently (i.e., one segment per RTT)
  - When cwnd **>=** twnd and queueing delay (σ) **>=** queueing budget
    - Reset cwnd to twnd and increase cwnd gently

XRC key behaviors

- XRC's behavior using an empty pipe between 0-5 secs
- XRC's behavior when competing with a flow that is slow in ramping up to its fair capacity-share between 5-15 secs
- XRC's behavior when competing with a flow that respects its fair-share of the capacity between 15-30 secs

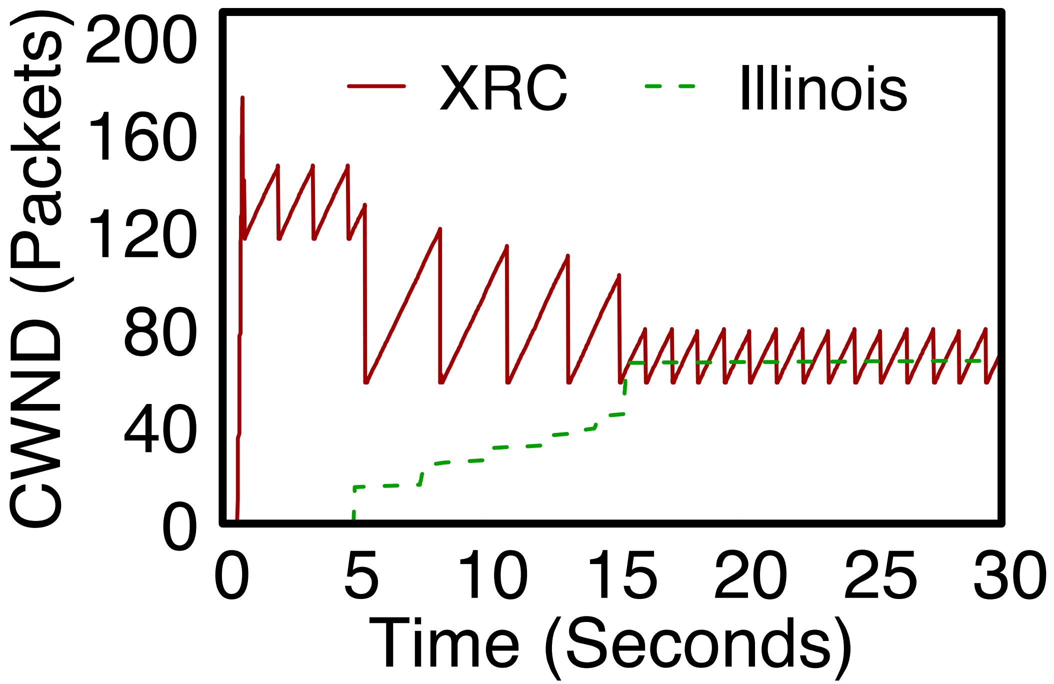

Evaluation

Fairness and Convergence

- Four XRC flows to arrive at the UE with an inter-arrival time of 10secs, with different RTTs. After 40secs, the flows leave the bottleneck with a 10secs gap between each flow's departure

XRC

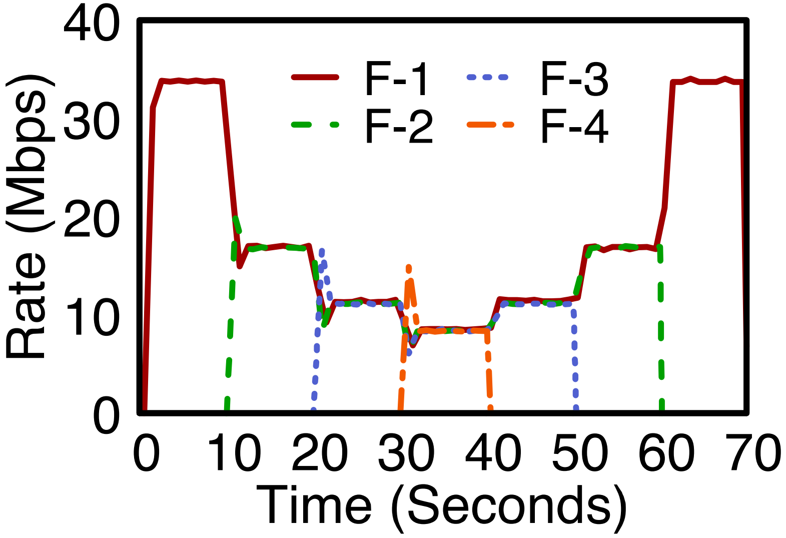

Jain's fairness index = 99.5%

CUBIC

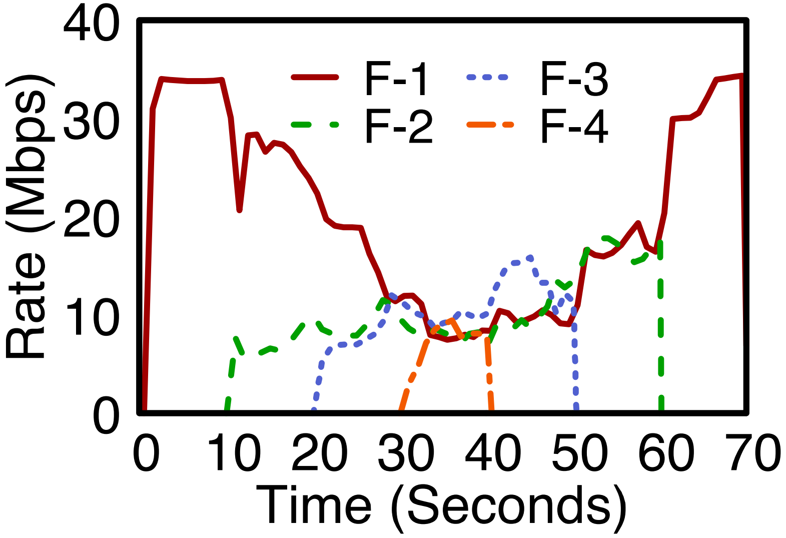

Jain's fairness index = 85.7%

Efficiency and Co-existence

- A long-lived TCP flow carries file download traffic which competes with a CUBIC flow carrying video content

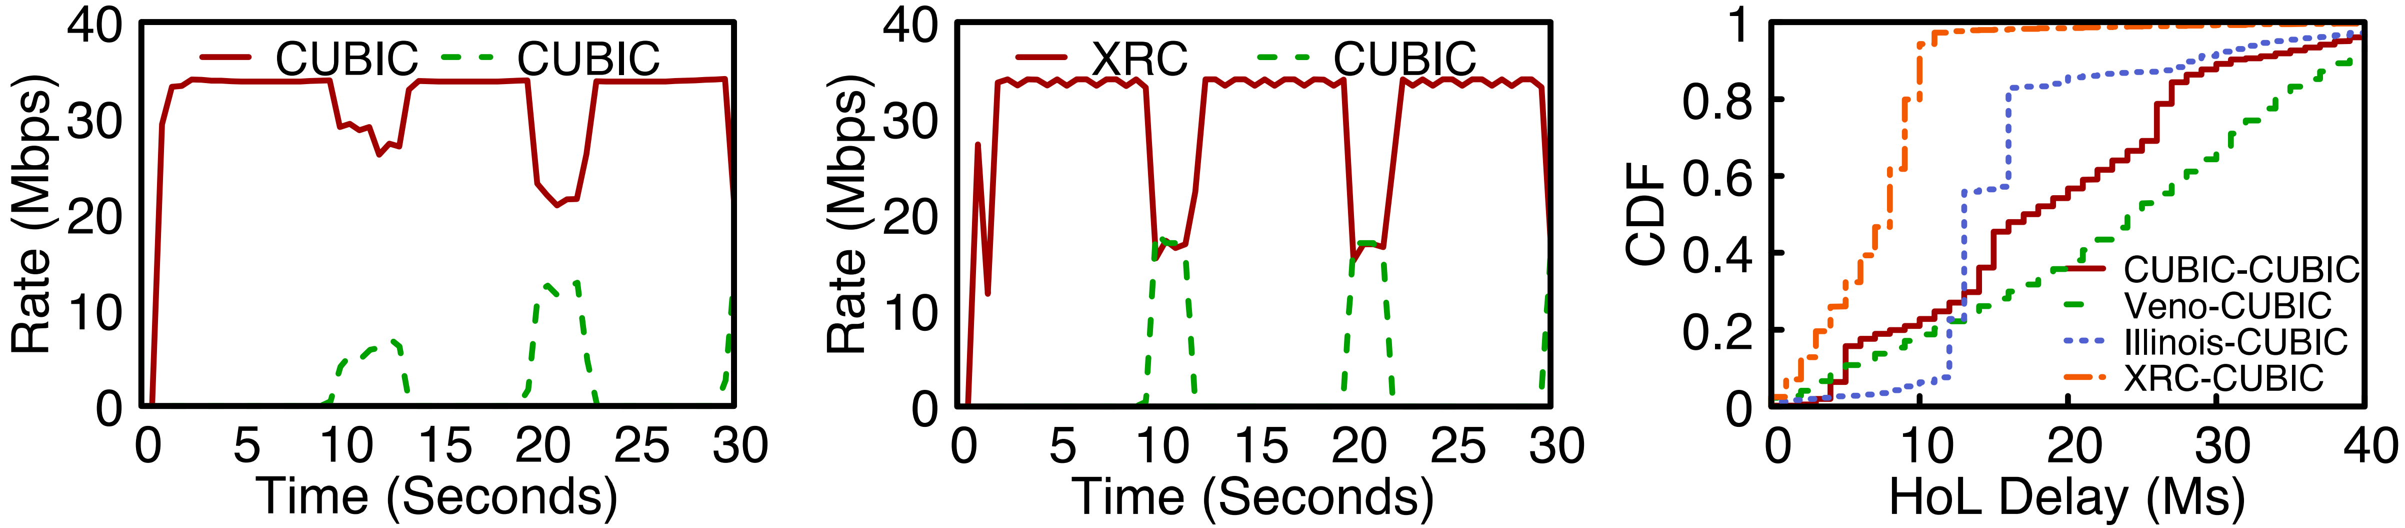

Existing TCP variants fail to fairly coexist when the competing flows are heterogeneous while XRC achieves fair co-existence
XRC achieve 4x lower queuing delay than other schemes

Future dimensions

- A real prototype with srsRAN and Linux Kernel supporting several state-of-the-art congestion control (e.g., BBR)
- XRC in the context of Multi-Path TCP (MPTCP)
- XRC in the context of the 3GPP's ATSSS (Access Traffic Steering, Switching, and Splitting) framework


interdigital
We invent the technologies that make life boundless.