

INTERNET-SCALE DISTRIBUTED SYSTEMS

BREAKOUT SESSION


Ramesh Sitaraman, UMass Amherst (lead)

NSF CSR/NETS JOINT PI MTG
NOVEMBER, 2019

Contributors

- Javad Ghaderi, Columbia
- Mohammad Hajiesmaili, UMass Amherst
- Stratis Ioannidis, Northeastern
- Zhenhua Liu, Stony Brook
- Faisal Nawab, UCSC
- Miao Pan, University of Houston
- Zubair Shafiq, University of Iowa
- Srinivas Shakkottai, Texas A&M

SCOPE: LARGE DISTRIBUTED NETWORKS



Examples:

- Content Delivery Networks.
- Cloud platforms

Scale:

- Quarter-a-million edge servers
- 1000+ data centers
- 100's of Tbps of traffic
- Trillions of transactions/day
- Billions of users/devices

Topic List

- New/Emerging distributed services
- Security and privacy
- Operations: managing the network at scale
- Sustainability
- Role of AI/ML
- Theoretical Foundations
- Realistic experimentation at scale
- Role of industry

New/Emerging Distributed Services


360 Video Delivery:

- Can we deliver the soccer world cup to 10's of millions of viewers in high-quality AR/VR in near real-time?
- Combination of low-latency (~20ms) and high bandwidth

Edge Analytics:

- The distributed edge for real-time data collection and analytics processing. E.g., IoT analytics.

Security and Privacy




Defending the edge.

*E.g., distributed WAF,
Client reputation*

Tradeoff between security/privacy and performance
Hardware support or light-weight mechanisms for
security/privacy needed


Operations: Managing complexity at scale



*Deployment, provisioning,
resource allocation,
monitoring, alerting, cost-
performance optimization*

- Distributed online algorithms for multi-criteria optimization
- ML for caching, load balancing, request routing, provisioning
- Explainability is important

Realistic Experimentation at Scale



- More NSF joint programs with industry
- NSF to facilitate data sharing from government and industry
- Create large distributed edge testbed for community use