BAM! Radio: A flexible software-defined radio platform for rapid prototyping of multi-hop wireless ad-hoc networks

Presenter: Dennis Ogbe, School of Electrical and Computer Engineering, Purdue University, West Lafayette, Indiana, USA

PURDUE UNIVERSITY | 150 YEARS OF GIANT LEAPS

Wireless networks as fertile ground for research

- Emerging networking technologies differ substantially from the current state of the art (LTE, Wi-Fi)

Internet of things: Low power, low throughput, massive random access

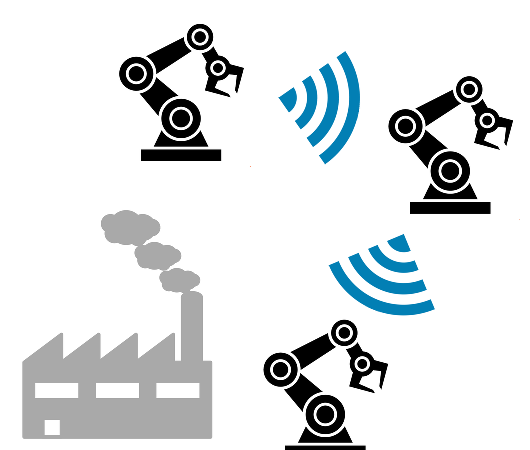

Machine-type comms.: Industrial applications, high reliability, low latency

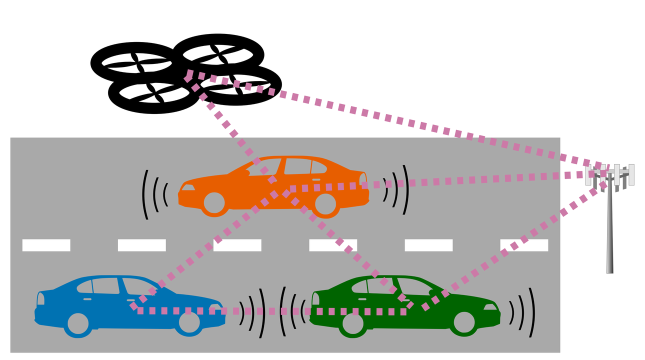

Vehicular networks: (autonomous) V2V/V2X, airborne vehicles, high mobility, rapid change in nodes

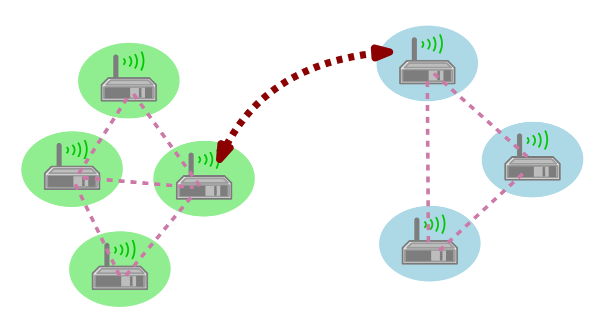

Collaborative networks: Advanced control & spectrum sharing mechanisms, AI-enabled

- Variety of use cases, topologies, semantics, and control mechanisms
- Lots of unanswered questions, ranging from fundamental limits to protocol design

Increased push toward experimental verification

- DoD (see DARPA's Colosseum) and NSF (see PAWR initiative) encourage prototyping and experimentation on SDR (software-defined radio) based testbeds
- Prototype implementation as first step toward commercialization
- Testbeds ranging from city-scale wireless to fully emulated
- Massive investments by both government & industry

Image source: https://www.spectrumcollaborationchallenge.com/about

SDR as a solid foundation

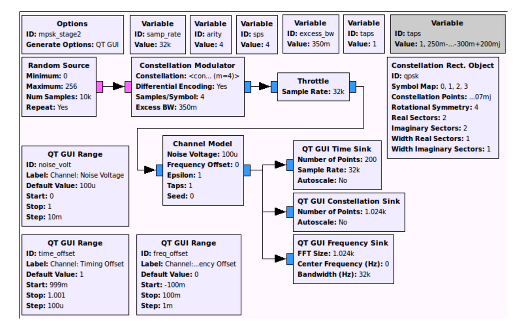

- SDRs like the Ettus USRP paired with libraries like GNU Radio provide a solid foundation for the implementation of prototype networks
- Problem: Focus either on a) ready-made applications for existing protocols (LTE, 802.11, ZigBee, etc.) or b) physical layer signal processing primitives (filtering, mixing, etc.)
- ***Need simple, extensible, multi-layer (PHY/DLL/NET) baseline radio network implementation to enable prototyping, experimentation, and research***

The BAM! Radio platform

- Developed as part of our entry into DARPA's Spectrum Collaboration Challenge
- IP-based packet radio stack (PHY/DLL/NET) + control mechanisms for ad-hoc networking
- Multiple physical layer implementations (currently: SC-FDMA for bulk data, 8-FSK for control links)
- Implemented in C++, extensible via Common Lisp (planned: Lua)
- Targets Ettus USRP x310, offloads some signal processing to NVIDIA GPUs
- Key: Designed to be extensible & hackable

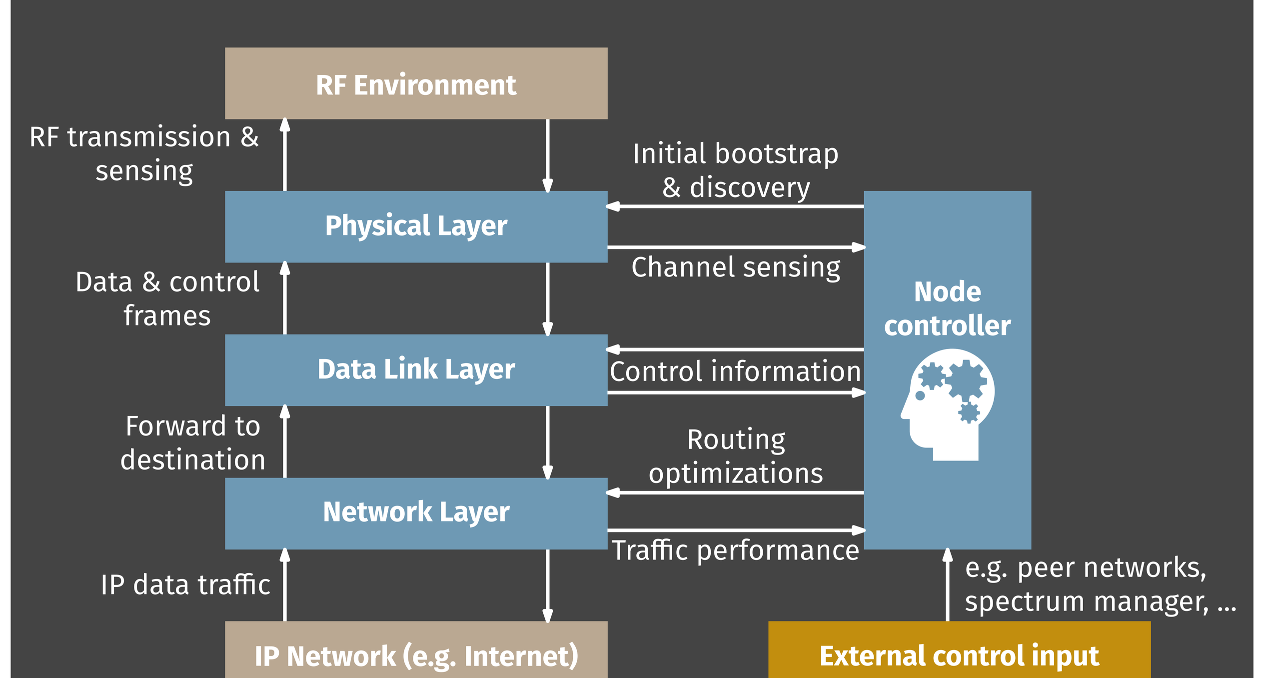

Image sources: https://www.ettus.com/all-products/x310-kit/, https://wiki.gnuradio.org/index.php/Main_Page

Example use case: DARPA Spectrum Collaboration Challenge

- DARPA's Grand Challenge to solve the "spectrum crunch" problem
- Goal: Channel recent advances in machine learning and artificial intelligence for spectrum management
- 3 years, 30+ teams, 13 finalists, 10 teams on stage at MWC LA 2019
- BAM! Wireless + BAM! Radio: Funded team for 3 years, phase 3 finalist, > $1.5M prize money
- Competition set-up: Submit code for one radio node, multiple instances are run as networks in "Colosseum"–A 128x128 MIMO channel emulator
- Different scenarios model real-world use cases (radar incumbent, disaster relief, etc.)

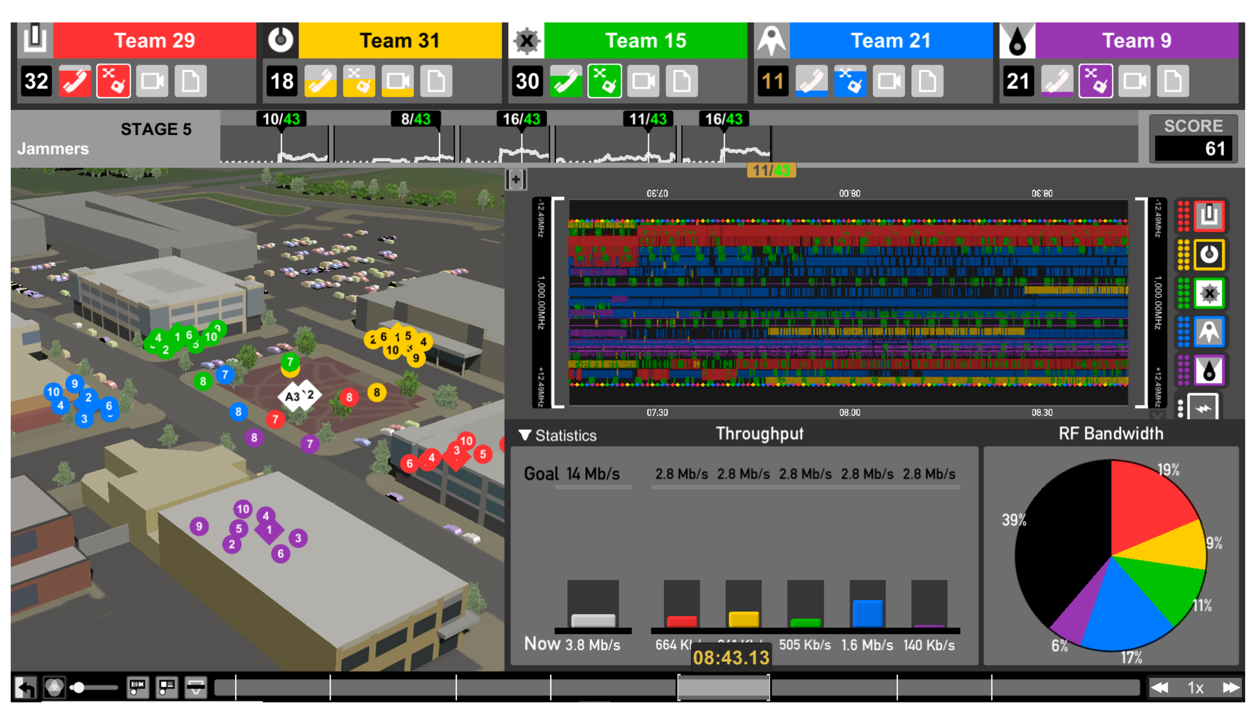

- Collaboration: Match score $\propto \min_{n \in 1,\dots,N} \mathrm{Score}_n$ (N = number of networks in scenario)
- Side channel between special nodes of each network

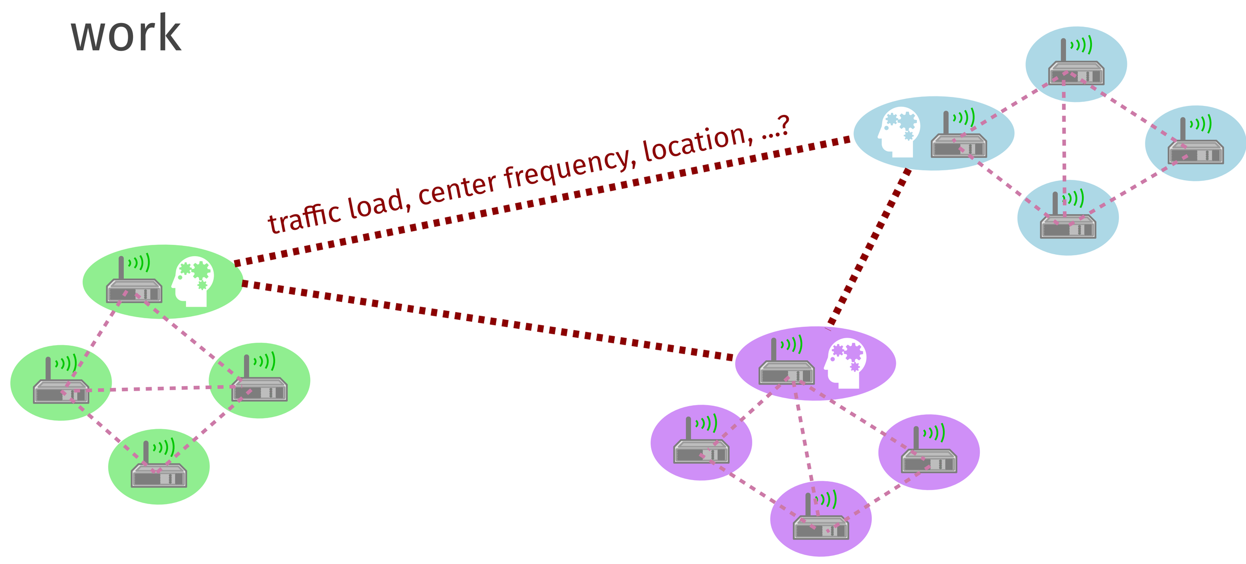

- Challenge: Protocol design, trusting external information, control/AI algorithms (currently done by Council of teams + DARPA)
- Many interesting research questions post-competition

Image source: https://www.spectrumcollaborationchallenge.com/media

Example use case: Low-latency multi-hop backhauling experiments

- The theoretical model of the separated multi-hop relay channel is central to many practical communication systems

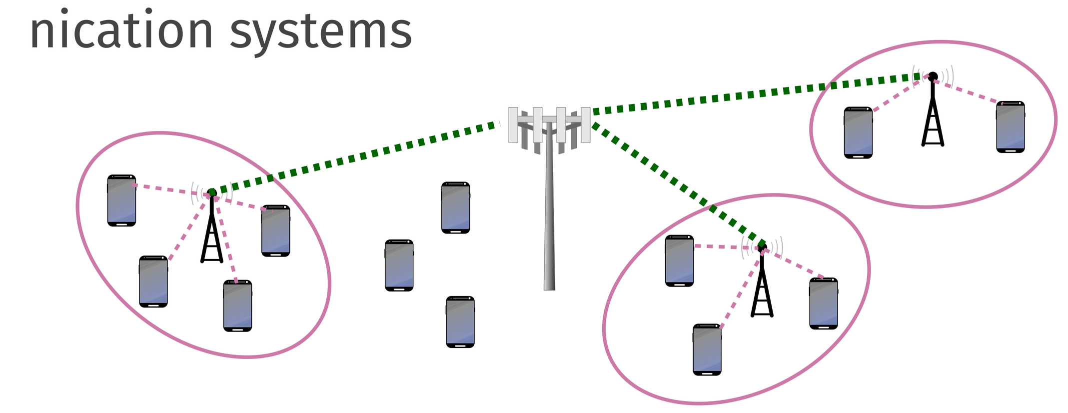

- Example: Small cells in 5G mmWave networks: Need wireless relays to obtain density required for mmWave

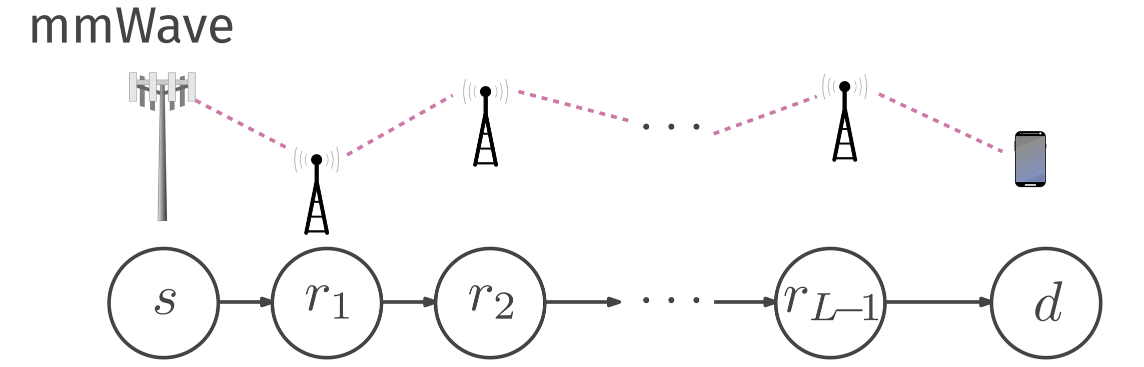

- Information-theoretic capacity analysis is clear (min-cut: $C = \min\left(\{C_l\}_{l=1}^{L}\right)$)
- For low-latency settings, the delay-throughput tradeoff is not well understood

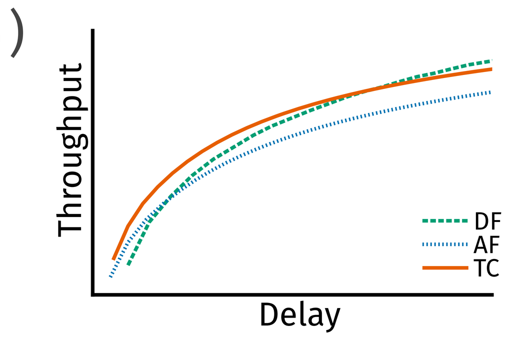

- Our work in [1] presented a simple coding scheme to improve on state-of-the-art (decode-&-forward) for short block lengths
- In [2] we propose a general analysis framework + more coding schemes for long block lengths

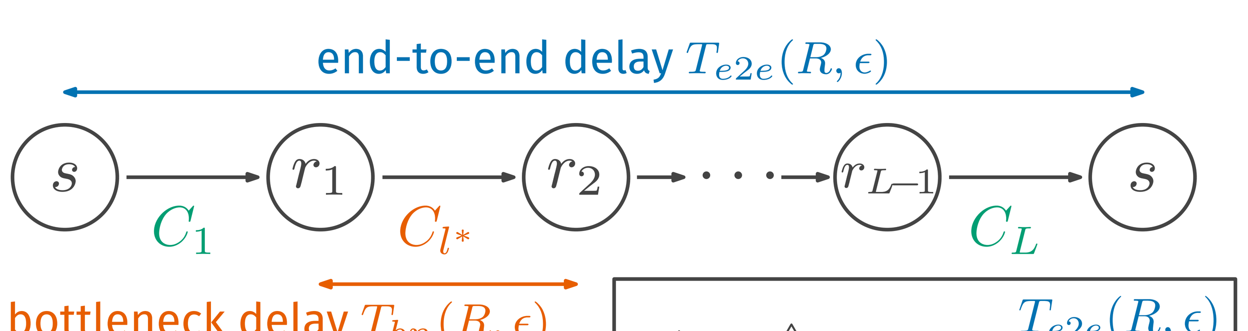

$$\mathrm{DAF}_\Phi \triangleq \lim_{R \nearrow C} \lim_{\epsilon \to 0} \frac{T_{e2e}(R,\epsilon)}{T_{bn}(R,\epsilon)}$$

- State of the art: $\mathrm{DAF}_{\mathrm{DF}} \approx \mathcal{O}(L)$
- Surprising results:
  (a) $\mathrm{DAF}_{\mathrm{TC}} = 1$ if $l^* = L$ (bottleneck is last)
  (b) $\mathrm{DAF}_{\mathrm{TC}} = 1$ for any l^* if feedback is allowed
- Future work: Use BAM! Radio to run real-world experiments based on this theory

Joint Work by D. Ogbe, C.-C. Wang, and D. J. Love

***Joint work by:** Dennis Ogbe†, Tomohiro Arakawa†, Stephen G. Larew†, Mai Zhang†, Andrew C. Marcum‡, David J. Love†, James V. Krogmeier†, Aly El Gamal†, Borja Peleato†, Nicolo Michelusi†, Chih-Chun Wang†, Alex Sprintson§
†Purdue University, ‡Raytheon BBN Technologies, §Texas A&M University

Selected Publications:
[1] C.-C. Wang, D. J. Love, and D. Ogbe, "Transcoding: A new strategy for relay channels," Allerton 2017, Oct. 2017.
[2] D. Ogbe, C.-C. Wang, and D. J. Love, "On the Optimal Delay Amplification Factor of Multi-Hop Relay Channels," ISIT 2019, July 2019.