

Compounding Effects of Medium Access and Congestion Control in Wireless Mesh Networks

Omer Gurewitz

Department of Communication Systems Engineering
Ben Gurion University, Israel

Joint work with:
Edward W. Knightly
and
Rice Networks Group

Project Goals

- Societal objective
 - Empower under-resourced communities through technology
- Technical objectives
 - Low-cost, high-performance, pervasive wireless
 - Minimal use of fiber
 - Applications providing new quality-of-life opportunities
 - Education and job-training
 - Work at home
 - Low cost health-care
 - Research platform: programmable and observable
 - Proof-of-concept for next generation protocols
 - Unique to demonstrate research advances "at scale"

Team

- Technology For All
 - Will Reed, Jim Forrest, and Nell Warnes
- University of Houston
 - Jerome Crowder
- Methodist Hospital Research Institute
 - Cliff Dacso
- Rice University
 - Rice Network Group

Geographical and Social Location

- Houston's East End
 - 37% of children below poverty
 - 56% have < \$25,000/year *household* income
 - Milby High School: 36-43% graduation rate
 - Health care: increased chronic and acute disease rates



Large-Scale Multi-hop Networks

- Technology For All Wireless Network
- Research platform: programmable and observable
- Wireless ISP for region since late 2004
- Over 4,000 users in 3 square kilometers
- Multi-tier architecture



Access Tier

- **Clients (homes and mobiles) to mesh nodes**
 - Clients (mobiles and residents) access mesh infrastructure
 - AP density: approximately 7 Mesh AP's per km²




Backhaul Tier

- Access nodes interconnected via backhaul tier
 - Access traffic forwarded to and from gateway
 - Omni directional 802.11 (b, g, or a)




Capacity Points

- Gateways inject capacity into backhaul tier
 - Injects capacity for sufficient Mb/sec/km²
 - Continued multi-hopping would render too many users over too many hops



TFA Network Architecture and Topology



- 802.11b access and backhaul serving 4,000 users
- 802.11a directional tier
- Multiple radios at gateway nodes, single radios elsewhere
- Opportunistic and population-driven GW locations
- 99% coverage over 4 km² when complete

Research Goals

- Passive monitoring of operational network with actual users
 - Understand user behavior and Channel condition
 - Compare with existing models
 - Modify existing models and simulators based on findings
 - Free access to traces
- Design new protocols and models
 - Channel conditions
 - Asymmetries
 - Topologies
 - Traffic matrices
- Proof of concept and verification
 - Validate protocols and model on operational network

Are you competing with the industry? i.e., what does it have to do with Research?



Case Study example

Compounding Effects of Medium Access and Congestion Control

Origins of starvation



- Demonstrate analytically as well as experimentally that the basic scenario of any mesh network is sufficient to induce starvation

Counter-starvation policy

- Suggest a simple solution which is supported by standard protocols

Prior Work: Performance Modeling

- Developed throughput model for arbitrary multi-hop topologies that predicts starvation
 - Captures inherent information asymmetries in multi-hop topologies
 - Role of asymmetric channel state inference



M. Garetto, T. Salonidis, and E. Knightly, "Modeling Per-flow Throughput And Capturing Starvation in CSMA Multi-hop Wireless Networks," in Proceedings of IEEE INFOCOM 2006, Barcelona, Spain, April 2006.

Prior Work: Congestion Control over Multi-Hop CSMA Networks

- It has been shown that TCP magnifies MAC unfair contention
 - Example: S. Lee K. Xu, S. Bae and M. Gerla, "TCP behavior across multihop wireless networks and the wired Internet," WOWMOM 2002.
- Solutions were suggested
 - Z. Fu, P. Zerfos, H. Luo, S. Lu, L. Zhang, and M. Gerla
 - INFOCOM'03 – LRED, Fixed window based on the hop count
 - K. Sundaresan, V. Anantharaman, H. Hsieh, and R. Sivakumar
 - MobiHoc'03 – ATP
 - T. Jimenez and E. Altman
 - PWC'03 – delayed TCP ack



However...


- Our prior understanding of "why TCP over multi-hop CSMA fails" is incorrect and has yielded solutions that are not effective
- It is believed that TCP pacing/smart dropping with the optimal pacing rate yields fairness but
 - we will show that starvation arises even with fixed sliding-window flow control coupled with CSMA (including 1!)
 - we will show that individual or aggregate one-hop flows can starve multi-hop flows

I went through the book but I am still starving

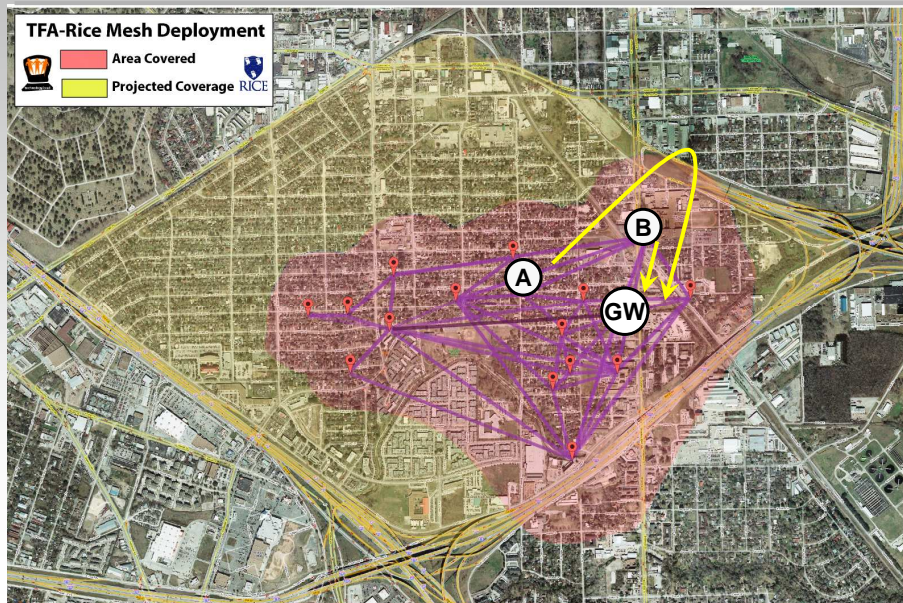


Basic Scenario of Mesh Network

- At least one Mesh Point for packet relay
 - Three node network
 - Node GW - Wired gateway
 - Node A and Node B – Mesh Points
- At least two TCP flows






Basic Topology in TFA



Severe Throughput Imbalance

Experiment of potential for starvation in operational mesh networks

- inject traffic from A and B to GW
- saturation conditions




The two-hop node "starves" when contending with the one-hop node

Origins of starvation

Compounding effect of three factors:

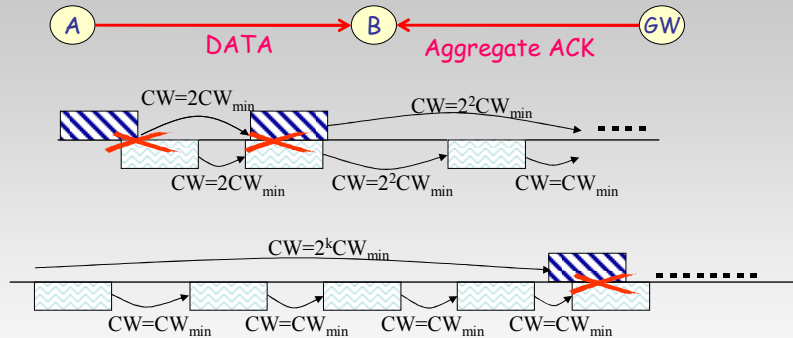
- Collision avoidance* in **medium access** protocol induces bi-stability in which pairs of nodes *symmetrically* alternate in capturing system resources
- Congestion control* in **transport protocol** induces *asymmetry* in the time spent in each state and favors the one-hop flow
- High penalty due to **cross-layer** effects in terms of loss, delay, and consequently, throughput, in order to re-capture system resources



Texas-size Starvation

Origins of starvation

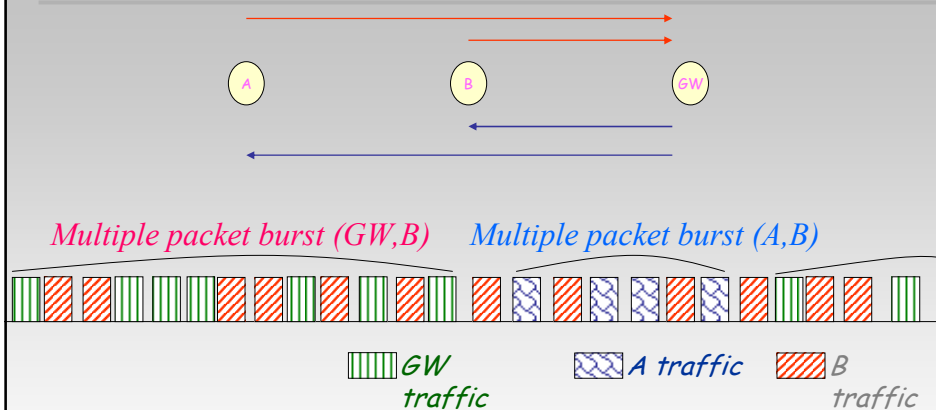
I) Medium Access and MAC Bi-stability



Due to lack of coordination:

- Bi-stable state: either A transmits and GW is in high backoff, or GW transmits and A is in high backoff
- *Success* state and *fail* state alternate
- *Symmetric behavior*

Middle Node Shares with Winner




- B is in range of both A and GW (complete channel state)
- B's packets interleave with A's and GW's packets

Origins of starvation


II) Asymmetry Induced by TCP

- Two nested transport loops and sliding windows



- *Asymmetric impact* of multipacket capture: transport loops change the duration of states

- (A, B) burst: the burst size is limited by:
 - TCP window size



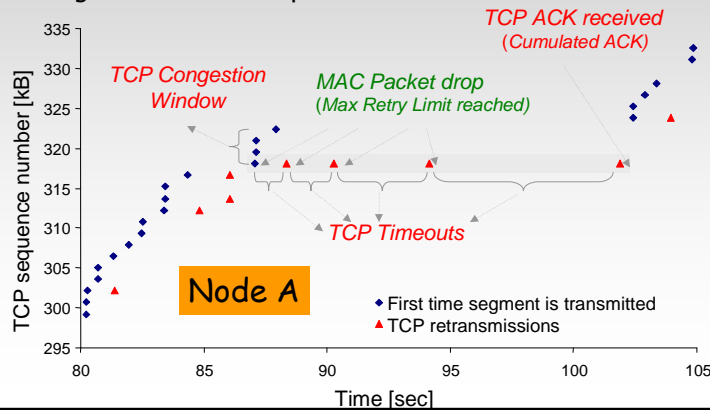
- (GW, B) burst: self-sustaining loop:
 - TCP ACK are generated



Origins of starvation

III) Severe Penalties

- Asymmetric impact of multipacket capture
 - Node GW incurs small penalty: short duration of fail state *but* long packet bursts
 - Node A incurs high penalty: long duration of fail state *and* low offered load high backoff & multiple TCP timeouts



Of course you are starving; the sniffer ate all your packets





Analytical Model

- Objectives
 - Isolate and capture the root cause of starvation
 - Model one aspect of congestion control (sliding window), queues, and CSMA/CA
- Technique
 - Embedded Markov chain model
 - Queue state, congestion window, contention window, carrier sense

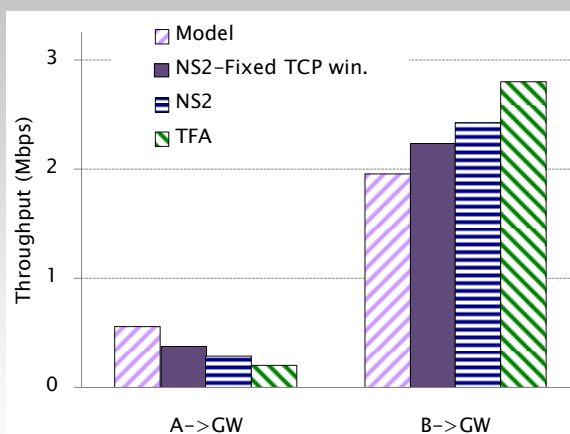
It is easy to see that the throughput is exactly:

$$\frac{\int_0^{\infty} \{4\sqrt{\alpha} + 17\omega^{2.5\alpha}\} \mu \alpha}{\int_0^{\infty} \pi x^{0.6\alpha} \sqrt{y^2 x dx dy}}$$



Evaluation: Model, Simulations, and TFA

- Model
 - static sliding window congestion control
- Simulation
 - fixed TCP congestion window + timeouts, cumulative ACKS, ...
 - legacy TCP New Reno (dynamic congestion window)
- Measurements at TFA
 - TCP New Reno+802.11



Model predicts starvation:
experimental factors exacerbate

What Next?

- Starvation exists
- Understand origins through analysis and modeling of protocols
 - CSMA + sliding window are sufficient
- Can we fix it?

*If it is not broken,
why mess with it?*



Counter-Starvation Policy

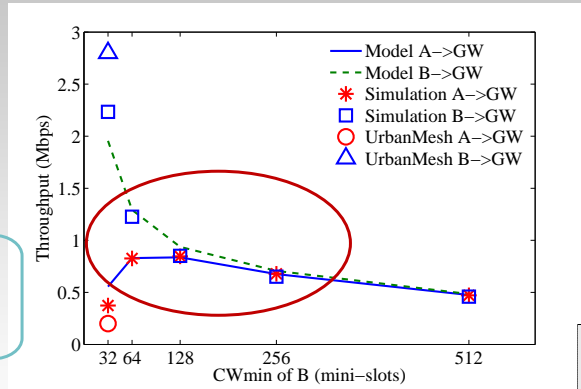
All nodes that are directly connected to the gateway should increase their minimum contention window to a value significantly greater than that of all other nodes

- Simple to implement- no overhead or message exchange between nodes
- Compliant with IEEE 802.11e EDCA
- Exploits congestion-control and MAC behavior for known topology



Validation: Does it work?


Solution: increase first hop contention window



Looks Great, but does it REALLY work?


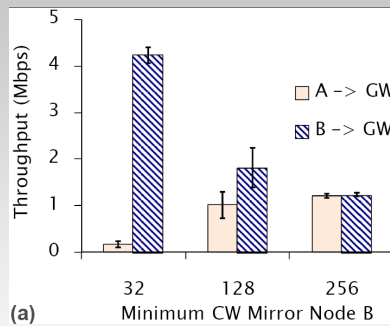


This is what I call a perfect match




Experimental Validation

RTS/CTS Enabled



Fairness can be achieved (under different definitions)



Broader Scenarios

Validate and evaluate the effect of the solution on more general scenarios:

- RTS/CTS disabled
 - Different packet size
 - Downstream
 - UDP traffic
 - Parking lot topology (longer chain topologies)
 - Multi-branch topologies
 - Multiple TCP flows
- Confirms that the starvation phenomenon exists in much broader scenarios
 - Proposed solution effective in more general topologies



Summary

- Analytically show that one-hop TCP flow(s) interacting with multi-hop TCP flow(s) is sufficient to induce starvation
- Demonstrate potential starvation in an operational multi-tier urban mesh network
- Analyze the joint effect of MAC and transport layer's congestion control on unfairness
- Suggest a simple Counter-Starvation Policy
- Implement and empirically validate the solution



Acknowledgements

- Sponsors



- Collaborators at Rice, TFA, and Methodist Hospital
- Students, postdocs, and staff
- More information:
 - <http://networks.rice.edu>

