Multi-Hop Wireless Networks:
Engineering Wireless Mesh Networks

Catherine Rosenberg,

This work was done in collaboration with Jun Luo and Andre Girard

University of Waterloo

Context: Scheduled Wireless Mesh Networks

- Fixed mesh routers form a **multi-hop wireless** network.
- Wireless ➔
  - Understanding & modelling physical layer is key (fading, **interference**, etc.).
  - Access scheme is key: **conflict-free scheduling (802.16, LTE)**.
- Primarily interested in **configuring** the network: routing, power and rate control, scheduling, gateway placement, etc.
- WLAN traffic is **aggregated** at the router and the traffic flows are (mostly) to or from the gateway**.**

Our approach: an offline configuration

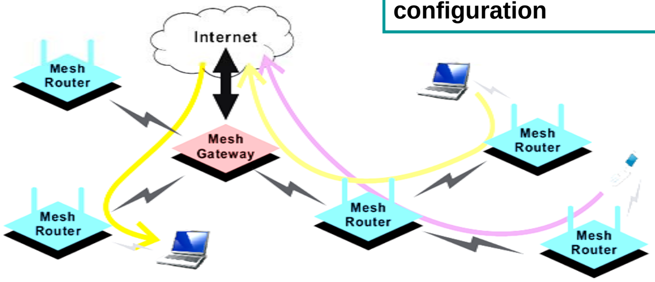

Questions:

- What can be gained by using multiple power levels and/or multiple modulation schemes?
- How more efficient is multi-path routing versus single path routing and how good is min-hop routing when compared with optimal routing?
- What is the relationship between spatial reuse and performance?
- What about lifetime?
- How important is the gateway placement?
- Can we revisit max-min throughput vs. proportional fairness?

Model

A. Karnik, A. Iyer and C. Rosenberg; "Throughput-optimal Configuration of Fixed Wireless Networks", in IEEE/ACM Transaction in Networking, vol 16, Issue 5, Oct. 2008, 1161-1174.

- **Max-min flow rate** – maximize the min end-to-end flow throughput achievable in the network (we also have results on Proportional Fair).
- **Cumulative interferences** from multiple links are considered: SINR (signal-to-interference-and-noise ratio) based model.
- **Multi-path routing**: An (o,d) flow can be split among a number of paths between o and d. These paths are given. Also single path routing.
- **Link scheduling**: **Central** control of link transmission. This is calculated as a fraction of the time an *Independent Set* (ISet) is allowed to transmit. An ISet is a set of links that can all transmit at the same time without creating harmful interference to each other.
- Nodes have access to multiple power levels and multiple modulations.
- Our **numerical** tools allow us to engineer **large networks**:
  - **input**: node/gateway locations, power levels, modulation schemes, channel gains, flows, objective function.
  - **output**: flow rates, routing, scheduling: which independent sets are used and for what proportion of time.

Assumptions and Problem Setup

- Traffic requirements are **static or quasi-static** (flow model):
- Channel gains are **almost time-invariant**:
  - Realistic in urban/suburban areas with roof-top antennas

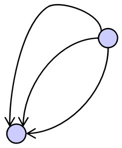

- Each link l is identified by
  - *o(l); d(l):* the origin and the destination nodes of *l*.
  - P_l: the tx power used by *o(l)*. It takes its value from a finite set ***P*** (i.e., power control ability).
  - c_l: the link rate in bits per second. It takes its value from a finite set ***C*** *(i.e.,* multi-rate ability).
- A link l is feasible if its *signal to noise ratio* $\mathrm{SNR}_l = \frac{G_l P_l}{N_0} \geq \beta(c_l)$ where G_l denotes the channel gain on *l,* N_0 is the average thermal noise power in the operating frequency band, and $\beta(c_l)$ is the threshold related to the modulation/coding scheme yielding c_l.
- The channel gain between two points separated by distance d is assumed to be given by $F_l(d/d_0)^{-\eta}$ where d_0 is the close-in reference distance, F_l is the shadowing and fading gain and η is the path loss exponent.

Additive Interference Model

- It is described using the concept of an *independent set (ISet):* A set of links that can all operate at the same time (i.e., the interference they produce is not harmful to any of the links in the set).
- First note that a set s of links is an ISet only if no two links in the set share a node.
- Second, for s to be an ISet it should meet

$$\gamma_l = \frac{G_l P_l}{N_0 + \sum_{l' \in s : l' \neq l} G_{l'l} P_{l'}} \geq \beta(c_l) \quad \forall\, l \in s.$$

where γ_l is the signal to interference plus noise ratio (SINR) of link l and $G_{l'l}$ is the channel gain from *o(l')* to *d(l)*.

Network Model

- Let F denote the set of **flows**. A flow f is identified by its source-destination pair $(f_s; f_d)$ and has a rate λ_f
- We focus in the following on the **multi-path routing** formulation.
- Denote by R_f the set of all routes that can be used by flow f and by R^l_f the set of all routes that can be used by f going through link l. The amount of flow f routed on r in R_f is Φ^r_f. Hence $\sum_{r\in\mathcal{R}_f}\phi^r_f=\lambda_f$ and $\mathbf{\Phi} = [\Phi^r_f]$ is the routing vector.
- **A link schedule** is a vector $\alpha=[\alpha_s]_{s\in\mathcal{I}}$ where $\mathcal{I}$ is the set of all ISets and such that $\alpha_s > 0$ if s is scheduled, otherwise $\alpha_s = 0$. We interpret α_s as the fraction of time allocated to a ISet s. Clearly: $\sum_{s\in\mathcal{I}}\alpha_s\le 1.$ By scheduling only ISets the schedule is **conflict-free**.
- Let **A** be the set of all conflict-free schedules.
- Link capacities under a conflict-free schedule α:

$$c_l(z_l,\alpha)=c_l z_l\sum_{i\in\mathcal{I}_l((z),(P))}\alpha_i$$

Formal Problem Statement

$$\max\lambda$$

$$\sum_{f\in\mathcal{F}}\lambda_f\left(\sum_{r\in\mathcal{R}^l_f}\phi^r_f\right)\le c_l(z_l,\alpha)\qquad l=1,\ldots,L$$

$$\sum_{r\in\mathcal{R}_f}\phi^r_f=1,\phi^r_f\ge 0\qquad f=1,\ldots,M$$

$$\lambda\le\lambda_f\qquad f=1,\ldots,M$$

$$\lambda\ge 0,\alpha\in\mathcal{A}(\mathbf{z},\mathbf{P}),\mathbf{z}\in\mathcal{Z}(\mathbf{P}),\mathbf{P}\in\mathcal{P}$$

- First set of constraints – link capacity constraints
  - LHS – Traffic imposed on link
  - RHS – Link capacity under conflict-free schedule α
- Third set – to maximize the minimum
- It is a standard but very large linear programming (LP) problem.

Notes on Problem Formulation

- Our formulation is very powerful and allows for numerous scenarios.
  - A version exists for single path routing.
  - Another exists for proportional fairness (in that case the problem becomes non-linear).
  - Another one for scheduling alone (when routing is fixed).
- How large is very large? The variables are the α_s's.
- If N is the number of nodes, and we assume 1 power and 1 modulation, we have potentially up to (approximately) N^2 links (depending on the value of P) and hence something of the order of 2^{N^2} subsets that need to be checked to know if they are ISets.
- If we have more than 1 power level or 1 modulation, we increase the number of potential links and hence of potential ISets.

Development of Computational Tools

- We have developed a **column generation technique** which allows us to solve **exactly** medium-size problems (up to 50 nodes or smaller but with several power levels or rates).
- The difficulty is to solve the NP-hard pricing subproblem in an efficient manner. We do that by introducing a technique that we call "greedy pricing" which uses an enumeration-based solver on a restricted set of links.
- We have shown that this technique allows us to compute exact solutions for problems much larger than what was possible before and is relatively fast.
- We have also proposed and compared two approximate algorithms that are fast and very accurate. They can be used to compute solutions for much larger networks.

J. Luo, C. Rosenberg, and A. Girard, "Joint Scheduling, Routing, Power Control and Rate, Adaptation in Fixed Wireless Mesh Networks: Algorithms and Engineering", submitted to IEEE/ACM Transaction in Networking, March 09 (under revision).

Engineering Insights: Scenarios

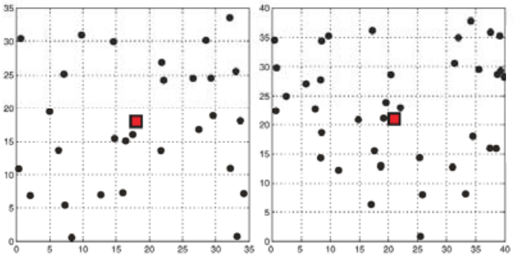

(a) Rand30a. (b) Rand40.

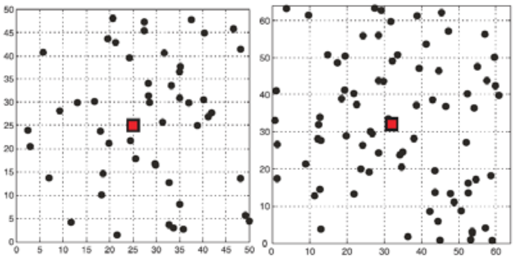

(c) Rand50a. (d) Rand80.

Modulation	coding rate	c_l	$\beta(c_l)$ (dB)
BPSK	1/2	1	6.4
QPSK	1/2	2	9.4
	3/4	3	11.2
16-QAM	1/2	4	16.4
	3/4	6	18.2

Link rates

d_0= 0.1m
F_l = 1 (no fading)
η = 3
N_0 = -100 dbm
Flows are converging by default, i.e., from all nodes to the gateway.

Early Results

- Assuming a single power level, the max-min throughput is a non-decreasing function of *P*. Hence, under this assumption, one should always use the highest possible transmit power if throughput is the only concern.

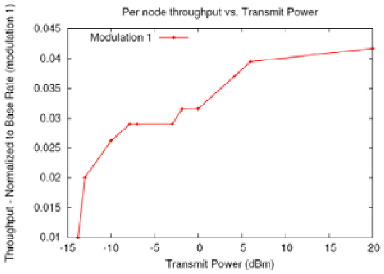

- The throughput is limited by the fact that the gateway cannot receive or transmit more than one packet at a time and hence is bounded by max(rate)/*N*.
- There are usually more than one optimal configurations. Usually these configurations are so complex that no simple rule can be deduced from them.

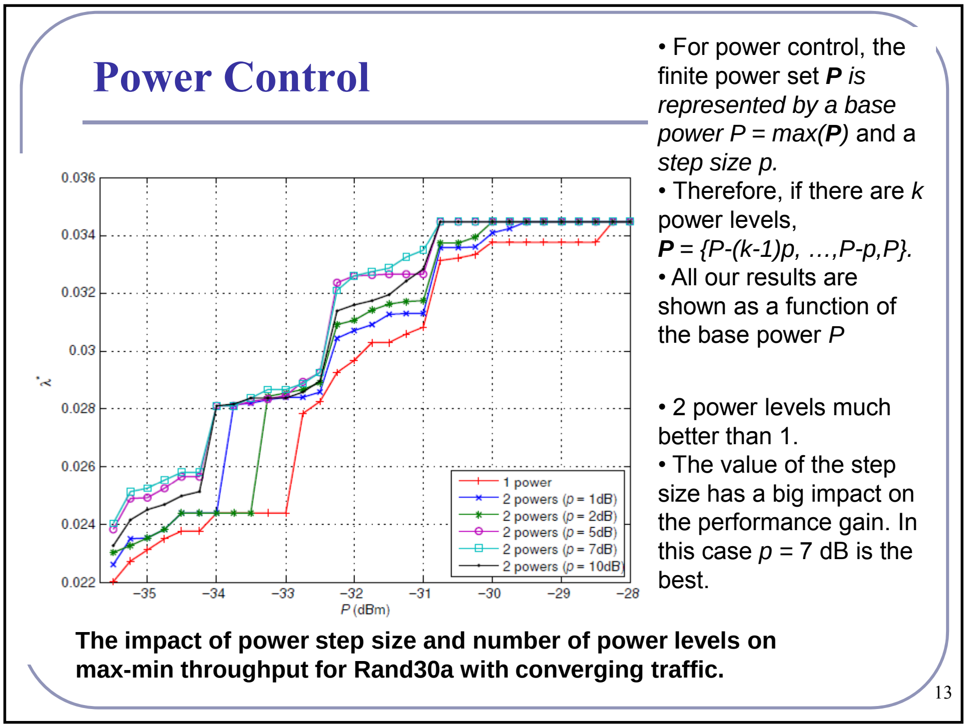

Power Control

- For power control, the finite power set **P** *is represented by a base power P = max(**P**)* and a *step size p.*
- Therefore, if there are *k* power levels, ***P*** *= {P-(k-1)p, …,P-p,P}.*
- All our results are shown as a function of the base power *P*

- 2 power levels much better than 1.
- The value of the step size has a big impact on the performance gain. In this case *p* = 7 dB is the best.

The impact of power step size and number of power levels on max-min throughput for Rand30a with converging traffic.

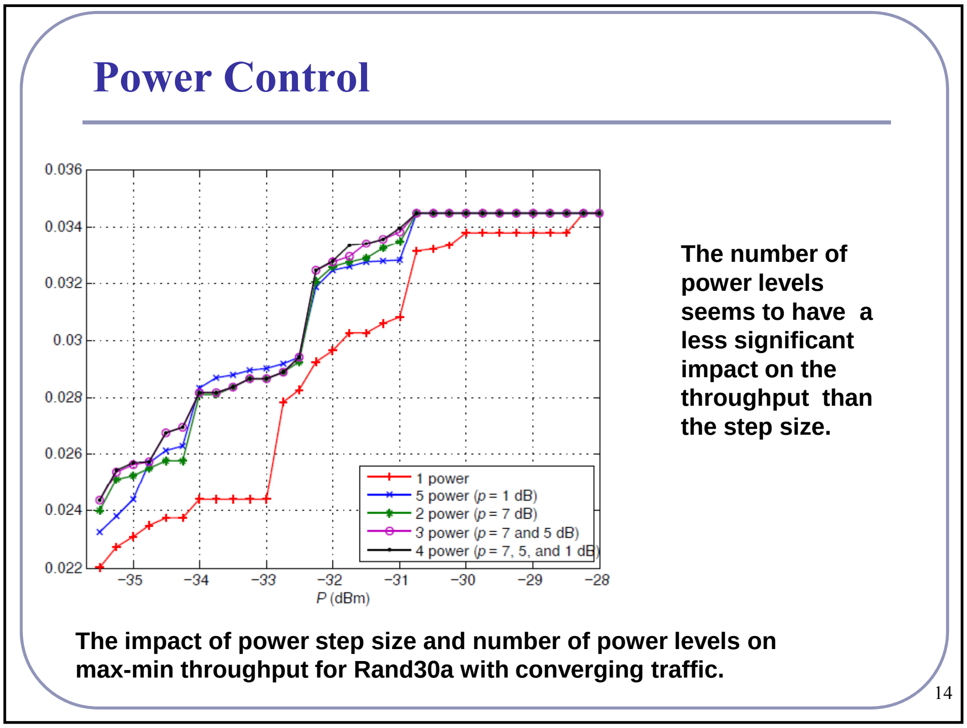

Power Control

The number of power levels seems to have a less significant impact on the throughput than the step size.

The impact of power step size and number of power levels on max-min throughput for Rand30a with converging traffic.

Rate Control

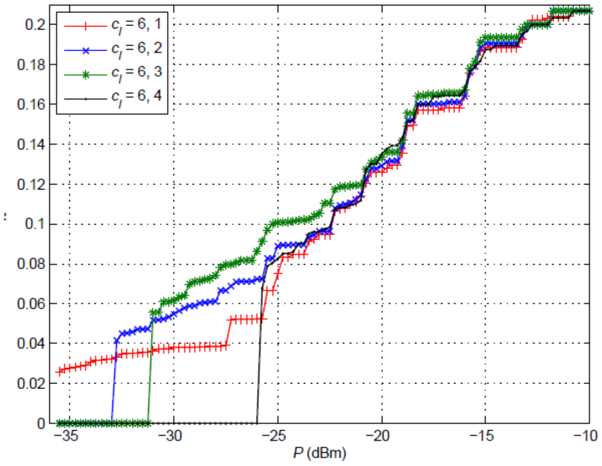

Existence of a "better" pair of rates, and we can determine this pair relatively quickly with our tools.

Multi-rate for Rand30a: the case of 2 rates and 1 power

Rate Control

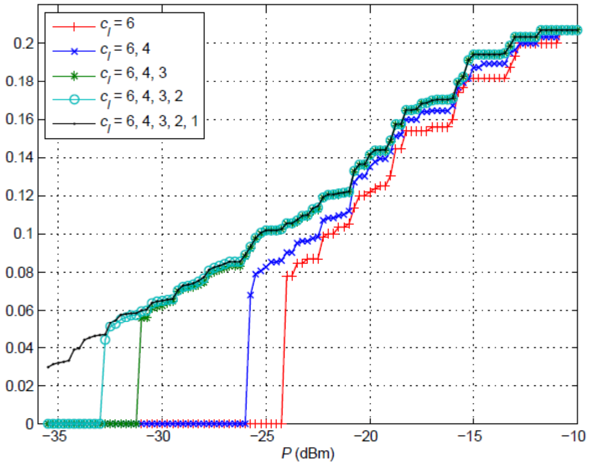

- Rate adaptation enables connectivity at lower power than a single high rate.
- Rate adaptation with 3 rates yields almost the same throughput than those with 4 and 5 rates after the network becomes connected.
- This seems to indicate that at an optimal configuration, only links with relatively high rates are used.

Multi-rate vs. single-rate for Rand30a

Multi-Path Advantage and Min-Hop Routing

- How much do we gain in throughput by allowing each flow to be routed on as many routes as necessary?
- Is min-hop routing a good routing scheme for scheduled mesh network?

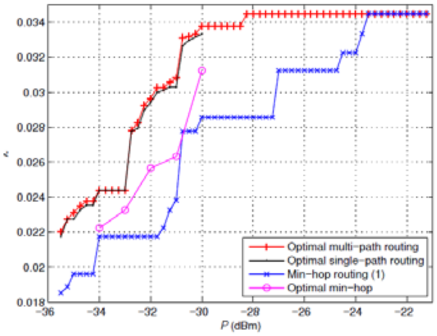

- Clearly, there is not much advantage in multi-path (optimal) routing
- Min-hop routing is far from optimal even if chosen optimally.

Rand30a with converging traffic, single power, single rate

Multihop Advantage (assuming one power level and a unit rate)

- In the case of diverging flows, let P_s be the gateway transmit power that allows every node to have a single-hop connection with the gateway. The maximum achievable max-min throughput is 1/N.
- Let P* be the minimum transmit power for which this maximum achievable throughput can be obtained via multi-hopping.
- The multi-hop advantage is P_s/P^*. It indicates how much the transmission power at the gateway can be decreased thanks to multihop without affecting the network performance.
- Multi-hop networking achieves the maximum achievable max-min throughput with a transmit power at the gateway often 4 or more times lower than the power needed for single-hop communication.
- This is made possible by allowing spatial reuse.

Network	P_s (dBm)	P* (dBm)	P_s/P^* (dB)
Grid25 (diverging)	-13.75	-18.74	5.00
Grid36 (diverging)	-11.00	-15.50	4.50
Rand30a (diverging)	−22.50	−29.00	6.50
Rand30b (diverging)	−23.25	−28.75	5.50
Rand50a (diverging)	−18.25	−25.50	7.25
Rand50b (diverging)	−19.00	−27.25	8.25

Revisiting Spatial Reuse

- A common belief is that the advantage of multi-hopping stems from spatial reuse and that the more spatial reuse the better. This is related to the size of the independent sets.
- The conjecture is that an optimal configuration would rely heavily on large ISets. In a 50 node network, there exist ISets of size up to 12.
- First we compute the optimal throughput curves without any restrictions on the size of the ISets.
- Then we compute the throughput obtained by restricting the size of the ISets that can be used to be less or equal to 1, 2, 3 and 4.

Revisiting Spatial Reuse

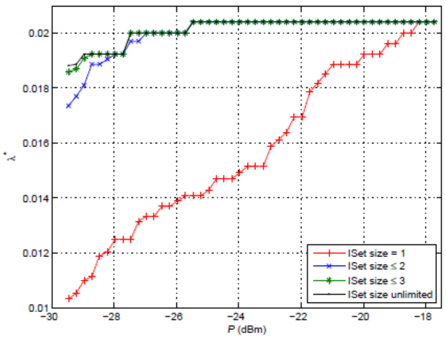

Rand50a (**diverging**), single power and rate.

- There is a big advantage to allow some level of spatial reuse.
- The gain obtained by allowing more spatial reuse, e.g., ISet ≤ 3 is not high as compared to using ISet ≤ 2
- The max-min throughput obtained by limiting the size of the ISets to 2 is never more than 10% below the optimal value.

Revisiting Spatial Reuse

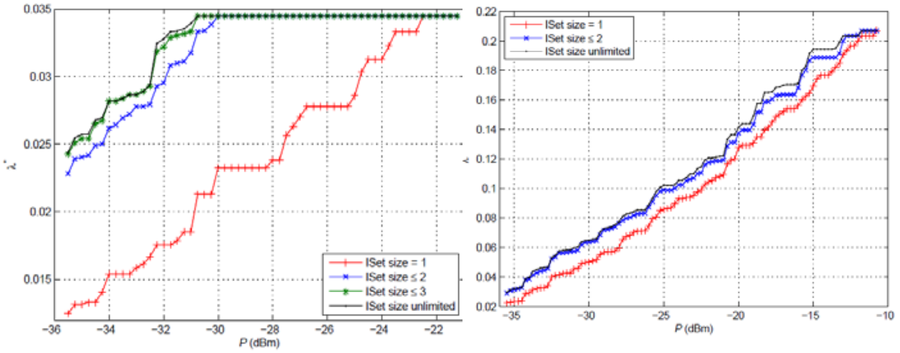

4 power levels

5 rates

Optimal max-min throughput under constraints on the maximum size of the ISets: Rand30a with converging traffic.

Single Gateway Placement

- Gateway placement is necessary in arbitrary networks and no placement is optimal for all Ps.

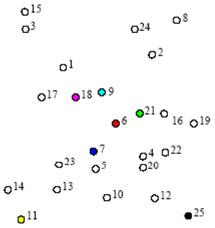

Arbitrary Network Topology

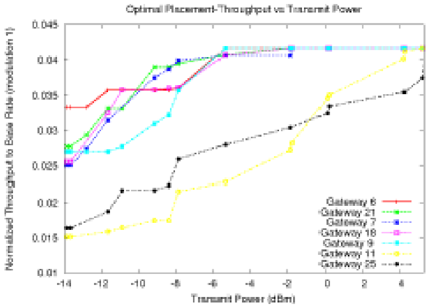

Optimal Throughput Curves

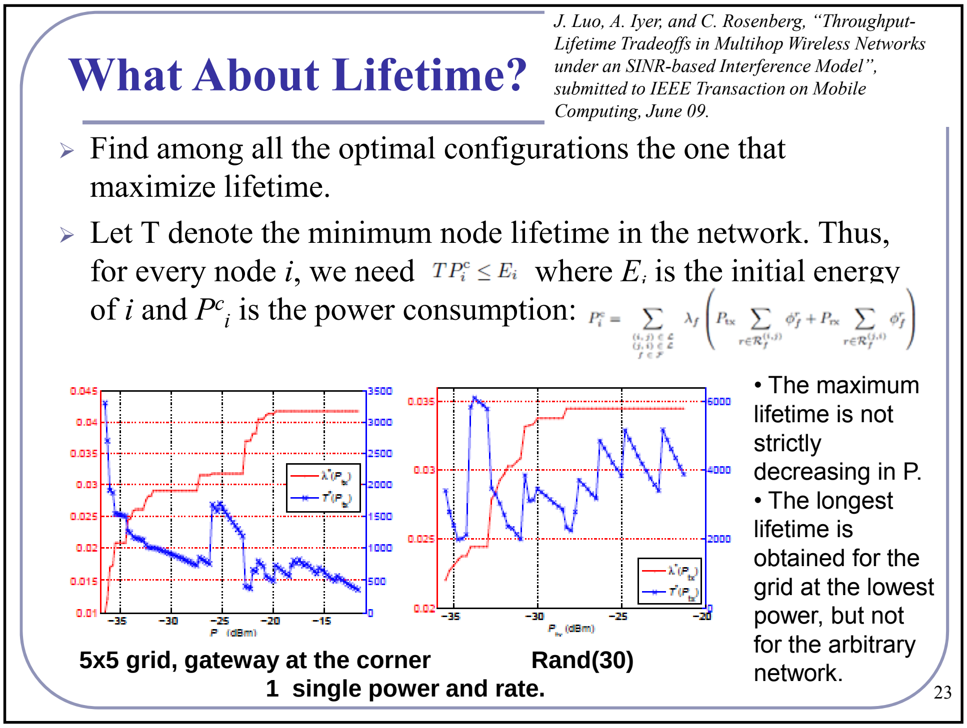

What About Lifetime?

J. Luo, A. Iyer, and C. Rosenberg, "Throughput-Lifetime Tradeoffs in Multihop Wireless Networks under an SINR-based Interference Model", submitted to IEEE Transaction on Mobile Computing, June 09.

- Find among all the optimal configurations the one that maximize lifetime.
- Let T denote the minimum node lifetime in the network. Thus, for every node *i*, we need $TP_i^c \le E_i$ where E_i is the initial energy of *i* and P^c_i is the power consumption: $P_i^c = \sum_{\substack{(i,j)\in\mathcal{L}\\(j,i)\in\mathcal{L}\\f\in\mathcal{F}}} \lambda_f \left(P_{tx} \sum_{r\in\mathcal{R}_f^{(i,j)}} \phi_f^r + P_{rx} \sum_{r\in\mathcal{R}_f^{(j,i)}} \phi_f^r \right)$

5x5 grid, gateway at the corner **Rand(30)**

1 single power and rate.

- The maximum lifetime is not strictly decreasing in P.
- The longest lifetime is obtained for the grid at the lowest power, but not for the arbitrary network.

23

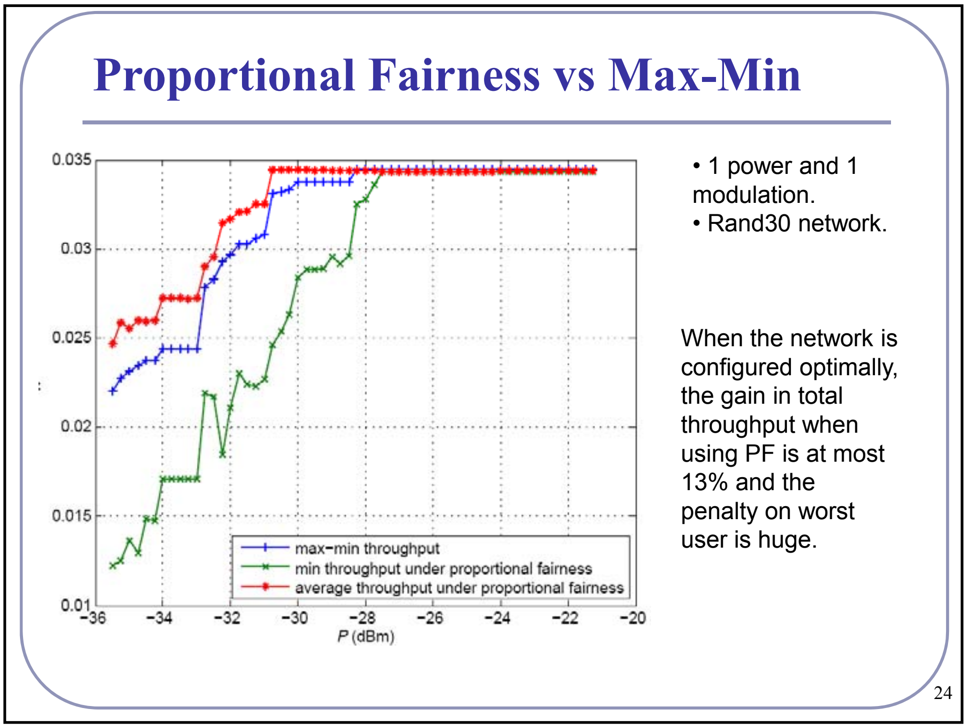

Proportional Fairness vs Max-Min

- 1 power and 1 modulation.
- Rand30 network.

When the network is configured optimally, the gain in total throughput when using PF is at most 13% and the penalty on worst user is huge.

24

Related Work

- Approaches to throughput maximization have (roughly) been of the following three kinds:
  1. offline design with exact solution (e.g., K. Jain, J. Padhye, V.N. Padmanabhan, and L. Qiu. *Impact of Interference on Multi-hop Wireless Network Performance*. In Proc. of the 9th ACM MobiCom, 2003.)
  2. offline design with approximate solution (e.g., G. Brar, D.M. Blough, and P. Santi. *Computationally Efficient Scheduling with the Physical Interference Model for Throughput Improvement in Wireless Mesh Networks*. In Proc of the 12th ACM MobiCom, 2006. P. Stuedi and G. Alonso. *Computing Throughput Capacity for Realistic Wireless Multihop Networks*. In Proc. of the 9th ACM MSWiM, 2006.),
  3. online dynamic control (e.g., L. Georgiadis, M.J. Neely, and L. Tassiulas. *Resource Allocation and Cross-Layer Control in Wireless Networks*. Foundations and Trends in Networking, 1(1):1–144, 2006. and the references therein).

Conclusions

- Our tools allow us to configure large networks.
- We obtain a lot of interesting engineering insights.

THANK YOU !