

Capacity

Chapter 11




ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Ad Hoc and Sensor Networks – Roger Wattenhofer – 11/1

Underwater Sensor Networks

- Static sensor nodes plus mobile robots
- Dually networked
 - optical point-to-point transmission at 300kb/s
 - acoustical broadcast communication at 300b/s, over hundreds of meters range.
- Project AMOUR [MIT, CSIRO]
- Experiments
 - ocean
 - rivers
 - lakes



Rating

- Area maturity

First steps

Text book

- Practical importance

No apps

Mission critical

- Theoretical importance

Not really

Must have

Ad Hoc and Sensor Networks – Roger Wattenhofer – 11/3

Overview

- Capacity and Related Issues
- Protocol vs. Physical Models
- Capacity in Random Network Topologies
- Achievable Rate of Sensor Networks
- Scheduling Arbitrary Networks

Ad Hoc and Sensor Networks – Roger Wattenhofer – 11/4

Fundamental Questions

- How much communication can you have in a wireless network?
- How long does it take to meet a given communication demand?
- How much spatial reuse is possible?
- What is the **capacity** of a wireless network?

- Many **modeling issues** are connected with these questions.
- You can ask these questions in many **different ways** that all make perfect sense, but give different answers.

- In the following, we look at a few results in this context, unfortunately only superficially.



Motivation

- Spatial capacity is an indicator of the “data intensity” in a transmission medium.

- The capacity of some well-known wireless technologies
 - IEEE 802.11b 1,000 bit/s/m²
 - Bluetooth 30,000 bit/s/m²
 - IEEE 802.11a 83,000 bit/s/m²
 - Ultra-wideband 1,000,000 bit/s/m²


- The wireless capacity is a function of the physical layer characteristics such as available **bandwidth or frequency**, but also how well the **protocols** on top of the physical layer are implemented, in particular media access. As such capacity is a theoretical framework for MAC protocols.



Protocol Model

- For lower layer protocols, a model needs to be specific about interference. A **simplest interference model** is an extension of the UDG. In the protocol model, a transmission by a node in at most distance 1 is received iff there is no conflicting transmission by a node in distance at most R, with $R \geq 1$, sometimes just $R = 2$.


- + Easy to explain
- Inherits all major drawbacks from the UDG model
- Does not easily allow for designing distributed algorithms/protocols
- Lots of interfering transmissions just outside the interference radius R do not sum up
- Can be extended with the same extensions as UDG, e.g. QUDG



Hop Interference (HI)

- An often-used interference model is hop-interference. Here a UDG is given. Two nodes can communicate directly iff they are adjacent, and if there is no concurrent sender in the k -hop neighborhood of the receiver (in the UDG). Sometimes $k = 2$.

- Special case of the protocol model, inheriting all its drawbacks
- + Simple
- + Allows for distributed algorithms
- A node can be close but not produce any interference (see picture)
- Can be extended with the same extensions as UDG, e.g. QUDG



Physical (SINR) Model

- We look at the signal-to-noise-plus-interference (SINR) ratio.
- Message arrives if SINR is larger than β at receiver

$$\frac{P_u}{d(u,v)^\alpha} \geq \beta \left(N + \sum_{w \in V \setminus \{u\}} \frac{P_w}{d(w,v)^\alpha} \right)$$

Power level of sender u

Path-loss exponent, $\alpha = 2, \dots, 6$

Noise

Distance between transmitter w and receiver v

Minimum signal-to-interference ratio, depending on quality of hardware, etc.

- Mind that the SINR model is far from perfect as well.

SINR Discussion

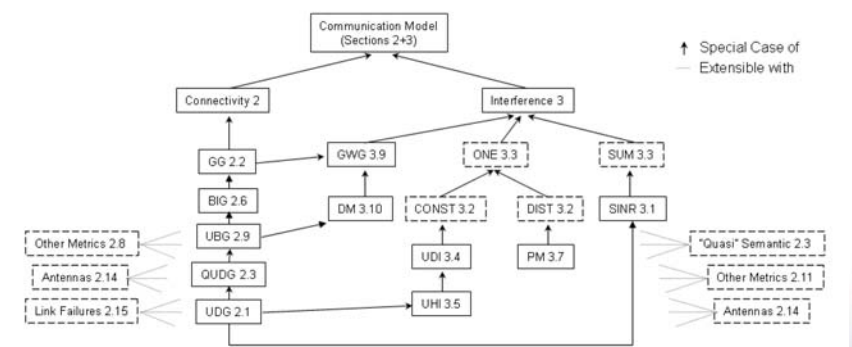
- In contrast to other low-layer models such as PM the SINR model allows for interference that does sum up. This is certainly closer to reality. However, SINR is not reality. In reality, e.g., competing transmissions may even cancel themselves, and produce less interference. In that sense the SINR model is pessimistic (interference summing up) and optimistic (if we remove the “1” from the SINR model, we have a UDG, which we know is not correct) at the same time.
- SINR is complicated, hard to analyze
- Similarly as PM, SINR **does not really allow for distributed algorithms**
- Also, in reality, e.g. the signal fluctuates over time. Some of these issues are captured by more complicated fading channel models.

More on SINR

- Often there is more than a single threshold β , that decides whether reception is possible or not. In many networks, a **higher S/N ratio** allows for more advanced modulation and coding techniques, allowing for **higher throughput** (e.g. Wireless LAN 802.11)
- However, even more is possible: For example, assume that a receiver is receiving two transmissions, transmission T_1 being much stronger than transmission T_2 . Then T_2 has a terrible S/N ratio. However, we might be able to “**subtract**” the strong T_1 from the total signal, and with $T - T_1 = T_2$, and hence also get T_2 .
- These are just two examples of how to get more than you expect.

Model Overview

[Algorithmic Models for Sensor Networks, Schmid et al., 2006]



- Try to prove **correctness** in an as “high” as possible model
- For **efficiency**, a more optimistic (“lower”) model is fine
- Lower bounds** should be proved in “low” models.

Measures for Network Capacity

- **Throughput capacity**
 - Number of successful packets delivered per time
 - Dependent on the traffic pattern
 - *E.g.: What is the maximum achievable, over all protocols, for a random node distribution and a random destination for each source?*
- **Transport capacity**
 - Network transports one **bit-meter** when one bit has been transported a distance of one meter
 - Number of bit-meters transported per second
 - *What is the maximum achievable, over all node locations, and all traffic patterns, and all protocols?*

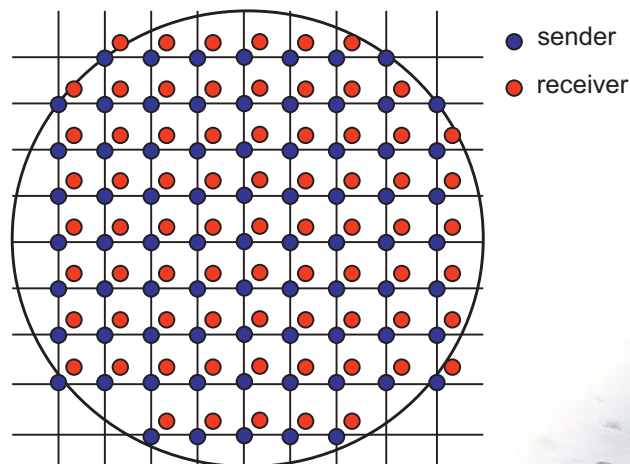


Transport Capacity

- n nodes are arbitrarily located in a unit disk
- We adopt the **protocol model** with $R=2$, that is a transmission is successful if and only if the sender is at least a factor 2 closer than any interfering transmitter. In other words, each node transmits with the same power, and transmissions are in synchronized slots.
- What configuration and traffic pattern will yield the highest transport capacity?
- Idea: Distribute $n/2$ senders uniformly in the unit disk. Place the $n/2$ receivers just close enough to senders so as to satisfy the threshold.




Transport Capacity: Example



Transport Capacity: Understanding the example

- Sender-receiver distance is $\Theta(1/\sqrt{n})$. Assuming channel bandwidth W [bits], transport capacity is $\Theta(W\sqrt{n})$ [bit-meter], or per node: $\Theta(W/\sqrt{n})$ [bit-meter]
- Can we do better by placing the source-destination pairs more carefully? Not really: Having a sender-receiver pair at distance d inhibits another receiver within distance up to $2d$ from the sender. In other words, it kills an area of $\Theta(d^2)$.
- We want to maximize n transmissions with distances d_1, d_2, \dots, d_n given that the total area is less than a unit disk. This is maximized if all $d_i = \Theta(1/\sqrt{n})$. So the example was asymptotically optimal.
 - BTW, a fun geometry problem: Given k circles with total area 1, can you always fit them in a circle with total area 2?



More capacities...

- The throughput capacity of an n node **random network** is $\Theta\left(\frac{W}{\sqrt{n \log n}}\right)$
- There exist constants c and c' such that $\lim_{n \rightarrow \infty} \Pr\left[c \frac{W}{\sqrt{n \log n}} \text{ is feasible}\right] = 1$
 $\lim_{n \rightarrow \infty} \Pr\left[c' \frac{W}{\sqrt{n \log n}} \text{ is feasible}\right] = 0$
- Transport capacity:
 - Per node transport capacity decreases with $\frac{1}{\sqrt{n}}$
 - Maximized when nodes transmit to neighbors
- Throughput capacity:
 - For random networks, decreases with $\frac{1}{\sqrt{n \log n}}$
 - Near-optimal when nodes transmit to neighbors
- In one sentence: **local communication is better...**




Even more capacities...

- Similar claims hold in the physical [SINR] model as well...
- Results are unchanged even if the channel can be broken into sub-channels
- There are literally thousands of results on capacity, e.g.
 - on random destinations
 - on power-law traffic patterns (probability to communicate to a closer node is higher)
 - communication through relays
 - communication in mobile networks
 - this research area has been quite a hobby of information theorists.
- Problem: The model assumptions are sometimes quite optimistic, if not unrealistic...
- What is the capacity of non-random networks?**



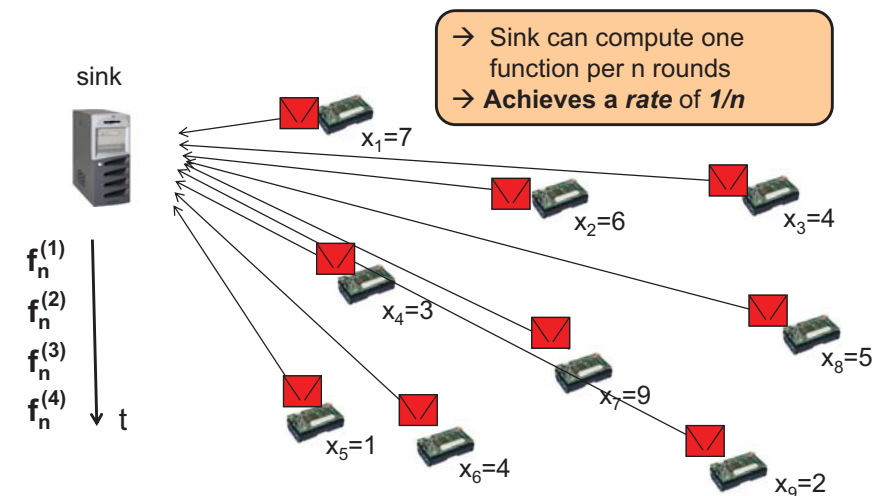
Data Gathering in Wireless Sensor Networks

- Data gathering & aggregation
 - Classic application of sensor networks
 - Sensor nodes periodically sense environment
 - Relevant information needs to be transmitted to **sink**
- Functional Capacity of Sensor Networks
 - Sink periodically wants to compute a **function** f_n of sensor data
 - At what **rate** can this function be computed?




The Simple Round-Robin Scheme

- Each sensor reports its results directly to the sink (one after another).



Is there a better scheme?!?



Capacity in Wireless Sensor Networks

At what **rate** can sensors transmit data to the sink?
 Scaling-laws \rightarrow how does rate decrease as n increases...?

$\Theta(1/n)$

$\Theta(1/\sqrt{n})$

$\Theta(1/\log n)$

$\Theta(1)$

Answer depends on

- Function to be computed
- Coding techniques
- **Network topology**
- **Wireless model**

Only simple functions
 (max, min, avg,...)
 No network coding



Practical relevance?

- Efficient data gathering!
- Efficient **MAC** layer!
- This (and related) problem is studied theoretically:


The Capacity of Wireless Networks
 Gupta, Kumar, 2000

- [Arpacioglu et al, IPSN'04]
- [Giridhar et al, JSAC'05]
- [Barrenechea et al, IPSN'04]
- [Liu et al, INFOCOM'03]
- [Grossglauser et al, INFOCOM'01]
- [Toumpis, TWC'03]
- [Gamal et al, INFOCOM'04]
- [Kysanur et al, MOBICOM'05]
- [Kodialam et al, MOBICOM'05]
- [Gastpar et al, INFOCOM'02]
- [Li et al, MOBICOM'01]
- [Mitra et al, IPSN'04]
- [Zhang et al, INFOCOM'05]
- [Bansal et al, INFOCOM'03]
- [Dousse et al, INFOCOM'04]
- [Yi et al, MOBIHOC'03]
- [Perevalov et al, INFOCOM'03]
- etc...

Network Topology?

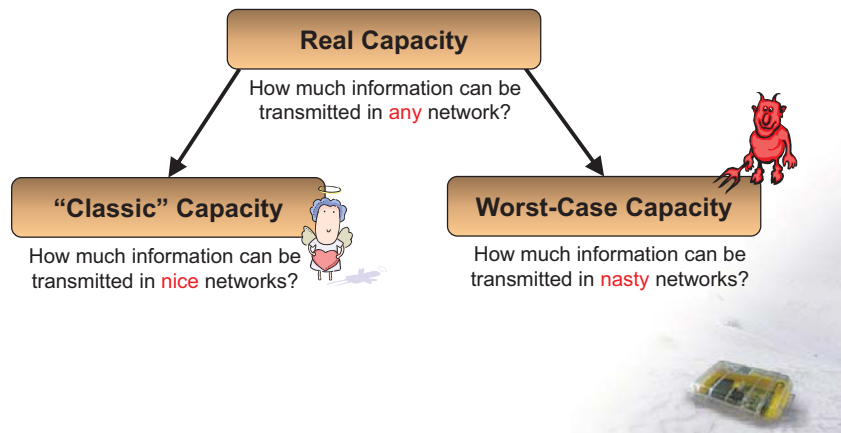
- Almost all capacity studies so far make very **strong assumptions** on node deployment, topologies
 - randomly, uniformly distributed nodes
 - nodes placed on a grid
 - etc.

What if a network looks differently...?



Capacity for Arbitrary/Worst-Case Network Topologies

- What can one say about worst-case node distributions?
- What can one say about arbitrary node distributions?




Wireless Models

- Several models for wireless communication
 - Connectivity-only models (e.g. UDG, QUDG, BIG, UBG, etc.)
 - **Interference models**
 - Protocol models
 - Two Radii model with constant power (e.g. UDG with interference radius $R=2$).
 - Nodes may use power control (transmission and interference disks of different size)
 - Physical models
 - SINR with constant power (every node transmitting with the same power)
 - SINR with power control (nodes can choose power)
 - Etc.
- Premise: **Fundamental results should not depend on model!**
 - And indeed, classical capacity (assuming e.g. random or regular node distribution) results are similar in all the models above
 - **Are there any examples where results depend on model?!**

Simple Example

A sends to D, B sends to C:



Assume a **single frequency** (and no fancy decoding techniques!)



Let $\alpha=3$, $\beta=3$, and $N=10nW$

Transmission powers: $P_B = -15$ dBm and $P_A = 1$ dBm


$$\text{SINR of A at D: } \frac{1.26mW/(7m)^3}{0.01\mu W + 31.6\mu W/(3m)^3} \approx 3.11 \geq \beta \quad \text{👍}$$

$$\text{SINR of B at C: } \frac{31.6\mu W/(1m)^3}{0.01\mu W + 1.26mW/(5m)^3} \approx 3.13 \geq \beta \quad \text{👍}$$

This works in practice!

[Moscibroda et al., Hotnets 2006]

- Measurements using mica2 nodes!
- Replaced standard MAC protocol by a (tailor-made) „SINR-MAC“
- Measured for instance the following deployment...




- Time for successfully transmitting 20,000 packets:

	Time required			Messages received	
	standard MAC	“SINR-MAC”		standard MAC	“SINR-MAC”
Node u_1	721s	267s	Node u_4	19999	19773
Node u_2	778s	268s	Node u_5	18784	18488
Node u_3	780s	270s	Node u_6	16519	19498

Speed-up is almost a factor 3

Worst-Case Rate in Sensor Network: Protocol Model

- Topology (worst-case!): Exponential node chain
- Model: **Protocol model, with power control**
 - Assume for simplicity that the interference radius is twice the transmission radius (however, this can be relaxed easily)



- Whenever a node transmits to another node all nodes to its left are in its interference range. In other words, no two nodes can transmit.
- Network **behaves like a single-hop network!**
- Same result for SINR with constant power or $P \sim d^\alpha$.

In the **protocol model**, the achievable rate is $\Theta(1/n)$.



Physical Model with Power Control

In the **physical model**, the achievable rate is $\Omega(1/\text{polylog } n)$, independent of the network topology.

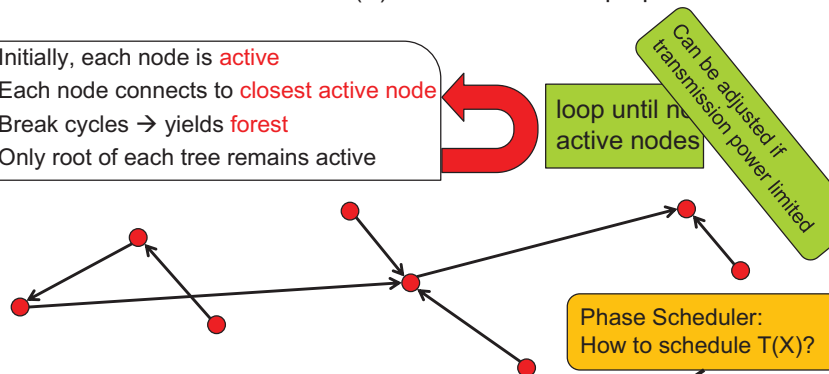
- Original result was $\Omega(1/\log^3 n)$. [Moscibroda et al, Infocom 2006]
- Later improved to $\Omega(1/\log^2 n)$. [Moscibroda, IPSN 2007]
- Algorithm is centralized, complex \rightarrow not practical
- But it shows that high rates are possible even in worst-case networks
- Basic idea: Enable **spatial reuse** by **exploiting SINR effects**.



Scheduling Algorithm – High Level Procedure

- High-level idea is simple
- Construct a hierarchical tree $T(X)$ that has desirable properties

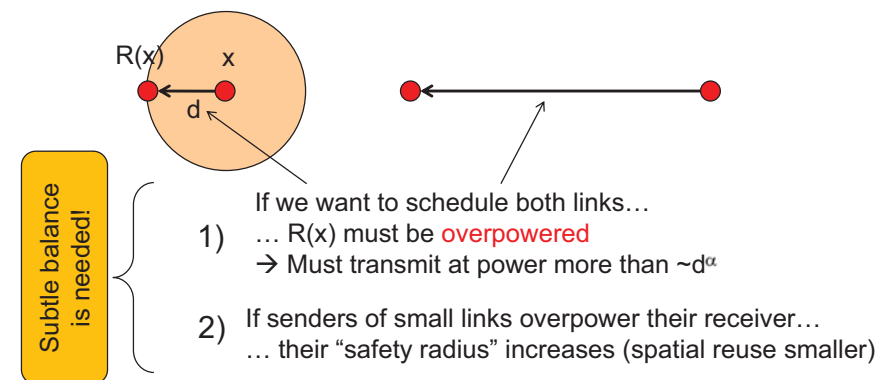
- 1) Initially, each node is **active**
- 2) Each node connects to **closest active node**
- 3) Break cycles \rightarrow yields **forest**
- 4) Only root of each tree remains active




The resulting structure has some **nice properties**
 \rightarrow If each link of $T(X)$ can be scheduled at least once in $L(X)$ time-slots
 \rightarrow Then, a rate of $1/L(X)$ can be achieved

Scheduling Algorithm – Phase Scheduler

- How to schedule $T(X)$ efficiently
- We need to **schedule links of different magnitude simultaneously!**
- Only possibility:
senders of small links must **overpower their receiver!**



Scheduling Algorithm – Phase Scheduler

- 1) Partition links into **sets** of similar length 
- 2) Group sets such that links a and b in two sets in the same group have at least $d_a \geq (\xi\beta)^{\xi(\tau_a-\tau_b)} \cdot d_b$
 - Each link gets a τ_{ij} value → Small links have large τ_{ij} and vice versa
 - Schedule links in these sets in one outer-loop iteration
 - Intuition: Schedule links of similar length or very different length
- 3) Schedule links in a group → Consider in **order of decreasing length** (I will not show details because of time constraints.)

Together with structure of $T(x) \rightarrow \Omega(1/\log^3 n)$ bound

Rate in Wireless Sensor Networks: Summary

		Worst-Case Capacity	Traditional Capacity
Power	Networks	Max. rate in arbitrary, worst-case deployment	Max. rate in random, uniform deployment
no power control		$\Theta(1/n)$	$\Theta(1/\log n)$
with power control		$\Omega(1/\log^2 n)$	$\Theta(1/\log n)$

[Giridhar, Kumar, 2005]

Exponential gap between protocol and physical model!

The Price of Worst-Case Node Placement


- Exponential in protocol model
- Polylogarithmic in physical model (almost no worst-case penalty!)

Theoretical Implications

- All MAC layer protocols we are aware of use either uniform or d^α power assignment.
 - Thus, the theoretical performance of **current MAC layer protocols** is in theory as bad as scheduling every single node individually!
- Faster polylogarithmic scheduling (faster MAC protocols) are theoretically possible in all (even **worst-case**) networks, if nodes choose their transmission power carefully.
 - Theoretically, there is **no fundamental scaling problem** with scheduling.
 - Theoretically efficient MAC protocols **must use non-trivial power levels!**
- Well, the word **theory** appeared in every line... ☺


Possible Applications – Improved “Channel Capacity”

- Consider a channel consisting of wireless sensor nodes
- What is the throughput-capacity of this channel...?




Possible Applications – Improved “Channel Capacity”

- A better strategy...
- Assume node can reach 3-hop neighbor



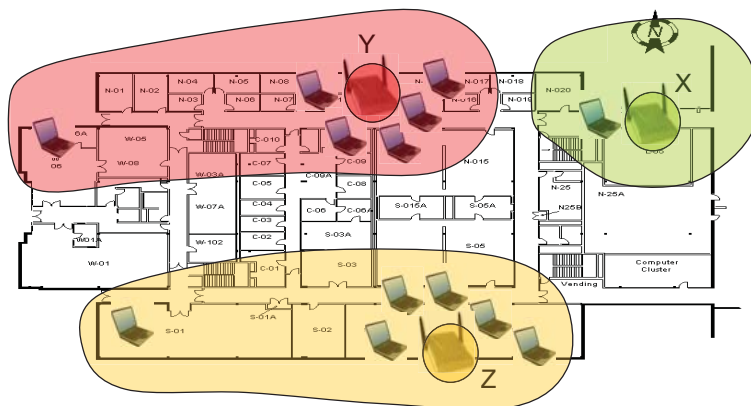
Possible Applications – Improved “Channel Capacity”

- All such (graph-based) strategies have capacity **strictly less than 1/2!**
- For certain α and β , the following strategy is better!



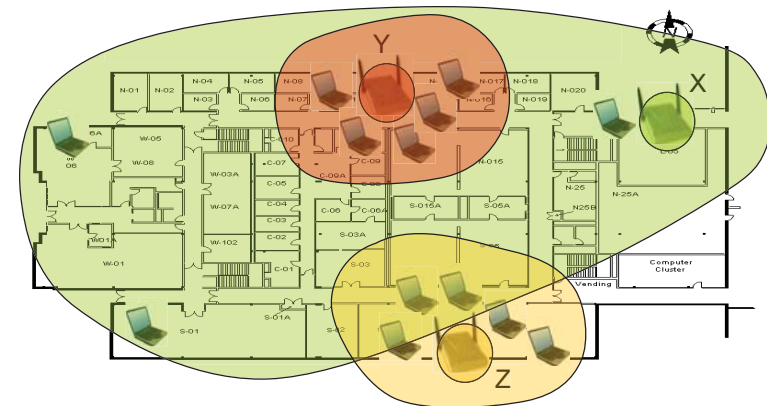
Possible Application – Hotspots in WLAN

- Traditionally: clients assigned to (more or less) closest access point
→ far-terminal problem → hotspots have less throughput




Possible Application – Hotspots in WLAN

- Potentially better: create hotspots with very high throughput
- Every client outside a hotspot is served by one base station
→ Better overall throughput – increase in capacity!



Scheduling Arbitrary Links?

- Given: A set of **arbitrary communication requests** in the plane
 - Each request is defined by position of source and destination
 - Each communication request has the same demand; if some request has a higher demand, just add links between the same sender/receiver
 - Just single-hop, no forwarding at intermediate nodes
 - Model: **SINR with constant power**
- Goal: **Minimize the time to schedule all links!**
 - Those scheduled in the same time slot must obey SINR constraints



- Example needs 3 time slots: **1,4,7** and **2,3,6** and **5,8**

Some Results

- Just checking whether some links can be scheduled in the same time slot is trivial. Simply test the SINR at each receiver.
 - In fact, even with power control this is easy. Since distances are fixed the SINR feasibility boils down to a set of linear equations:

$$\frac{P_i}{d_i^\alpha} \geq \beta \left(I + \frac{P_1}{d_1^\alpha} + \dots + \frac{P_{i-1}}{d_{i-1}^\alpha} + \frac{P_{i+1}}{d_{i+1}^\alpha} + \dots + \frac{P_k}{d_k^\alpha} \right), P_i \geq 0, \forall i.$$

- On the other hand, scheduling all links in minimum time is difficult (NP-complete), even with constant power. [Goussevskaia et al., 2007]
 - With power control, the complexity of scheduling is still unknown.
- What about **approximation algorithms**? Is it easy to schedule the maximum number of links in one slot? How much time do you need to deliver a given communication demand?

One-Slot Scheduling with Fixed Power Levels

- Given a set $L = \{l_1, \dots, l_n\}$ of arbitrary links, we want to maximize the number of links scheduled in one time-slot
- Constant approximation algorithm:

- Input: L ; Output: S ;
- Repeat
 - Add shortest link l_v in L to S ;
 - Delete all l_w in L , where $d_{wv} \leq c \cdot d_{vv}$;
 - Delete all l_w in L , where $a_S(l_w) \geq 2/3$;
- Until $L = \emptyset$;
- Return S ;

Schedule "strong" links first

Constant $c > 2$, $c = f(\alpha, \beta)$

Links with receivers s_w too close to sender r_v

Links with high affectedness*

- *Definition: Affectedness is how much interference a link can tolerate, i.e. $a_S(l_v) = 1$ if $\text{SINR}_S(l_v) = \beta$

Set of links S is valid iff $a_S(l_v) \leq 1$ for all l_v in S

One-Slot Scheduling: Correctness Proof

- We need to prove affectedness $a_S(l_v) \leq 1$ for all l_v in S


S_v^- : set of *shorter* (\leq) links in S , i.e., added *before* l_v

S_v^+ : set of *longer* (\geq) links in S , i.e., added *after* l_v

$a_{S_v^-}(l_v) \leq 2/3$ (OK! by algo)

$a_{S_v^+}(l_v) \leq 1/3$ (? See below!)

- All senders in set S_v^+ have pair wise distance $\delta = (c-1)d_{vv}$.
- We partition the space in infinitely many rings of thickness δ .
 - There is no sender in S_v^+ in circle R_0
 - A sender in ring R_k has at least distance $k\delta$
 - The number of senders in ring R_k is $O(k)$
 - Affectedness from ring R_k is $O(\beta k^{1-\alpha} \delta^{-\alpha})$
 - Total affectedness is $O(\sum_{k \geq 1} \beta k^{1-\alpha} \delta^{-\alpha}) \leq 1/3$, for $\alpha > 2$ and large enough constant
 - $c = f(\alpha, \beta)$



One-Slot Scheduling: Approximation Proof


- Count the number of links deleted by ALG that could have been scheduled in the optimum solution **OPT**: $OPT' = OPT \setminus ALG$

$$OPT' = OPT_1 + OPT_2$$

OPT_1 : links deleted in step 1

OPT_2 : links deleted in step 2

- Claim 1:** $|OPT_1| \leq \rho_1 |ALG|$, with $\rho_1 = f(c)$
- Proof: If the optimal wants to schedule more than ρ_1 links around receiver r_v , then two of these links have to be very close, and would not satisfy the SINR condition (since their length is at least the length of link l_v).

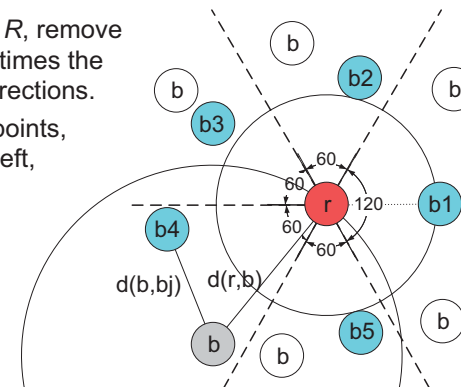


Helper Lemma: Blue-Dominant Centers Lemma

- Consider two disjoint sets B and R of *blue* and *red* points in the plane. If $|B| > 5q|R|$ then at least one point b in B is q -blue-dominant.

Definition: a point b in B is q -blue-dominant (q in \mathbb{Z}^+) if, for every ball of radius d around b , there are q times more blue than red points.

- Proof: for each red point r in R , remove q "guarding sets" of r , i.e., q times the 5 closest blue points in all directions.
- After processing all the red points, at least one blue point b^* is left, since $|B| > 5q|R|$. Point b^* is blue-dominant, since all red points in R are "guarded" by at least q blue points from all directions.



One-Slot Scheduling: Approximation Proof (Part 2)

- Claim 2:** $|OPT_2| \leq \rho_2 |ALG|$, $\rho_2 = 10$
 - Proof: Let $B = OPT_2$ (senders) and $R = ALG$ (senders also)
 - Assume for the sake of contradiction that $|B| > 10|R|$.
 - By the blue dominant centers lemma, there is a *2-blue-dominant* point b in OPT_2 . Since b is 2-blue-dominant, it is twice as much affected by OPT_2 than by ALG . (Really, it is the receiver of link l_b that is affected but thanks to the clearing of the neighborhood, the receiver is "close" to the sender; some nasty details omitted).
 - In other words, we have $a_{ALG}(l_b) < \frac{1}{2} \cdot a_{OPT_2}(l_b)$. With $a_{OPT_2}(l_b) \leq 1$ we have $a_{ALG}(l_b) < \frac{1}{2}$. This **contradicts** that link l_b has been deleted by step 2 of the algorithm, since step 2 only deletes links with $a_S(l_b) \geq \frac{2}{3}$.
- In summary, $(\rho_1 + \rho_2) \cdot |ALG| \geq |OPT_1| + |OPT_2| = |OPT'| = |OPT \setminus ALG| \geq |OPT| - |ALG|$. That is, $|ALG| \geq |OPT| / (\rho_1 + \rho_2 + 1) = \Omega(|OPT|)$.

Open problem

- This is an area with more open than closed problems. An obvious open problem is **scheduling with power control**. Formally, the problem can be defined as follows:
- A communication request consists of a source and a destination, which are arbitrary points in the Euclidean plane. Given n communication requests, assign a color (time slot) to each request. For all requests sharing the same color specify power levels such that each request can be handled correctly, i.e., the SINR condition is met at all destinations. The goal is to minimize the number of colors.
- Pretty much nothing is known about this problem.

