The Worst-Case Capacity of Wireless Networks

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Disclaimer…

- Work is about wireless networking in general
  - Presentation focusing on wireless sensor networks

- Joint Work
  - Thomas Moscibroda (thanks for some slides)
  - Olga Goussevskaia
  - Yvonne Anne Oswald
  - Yves Weber

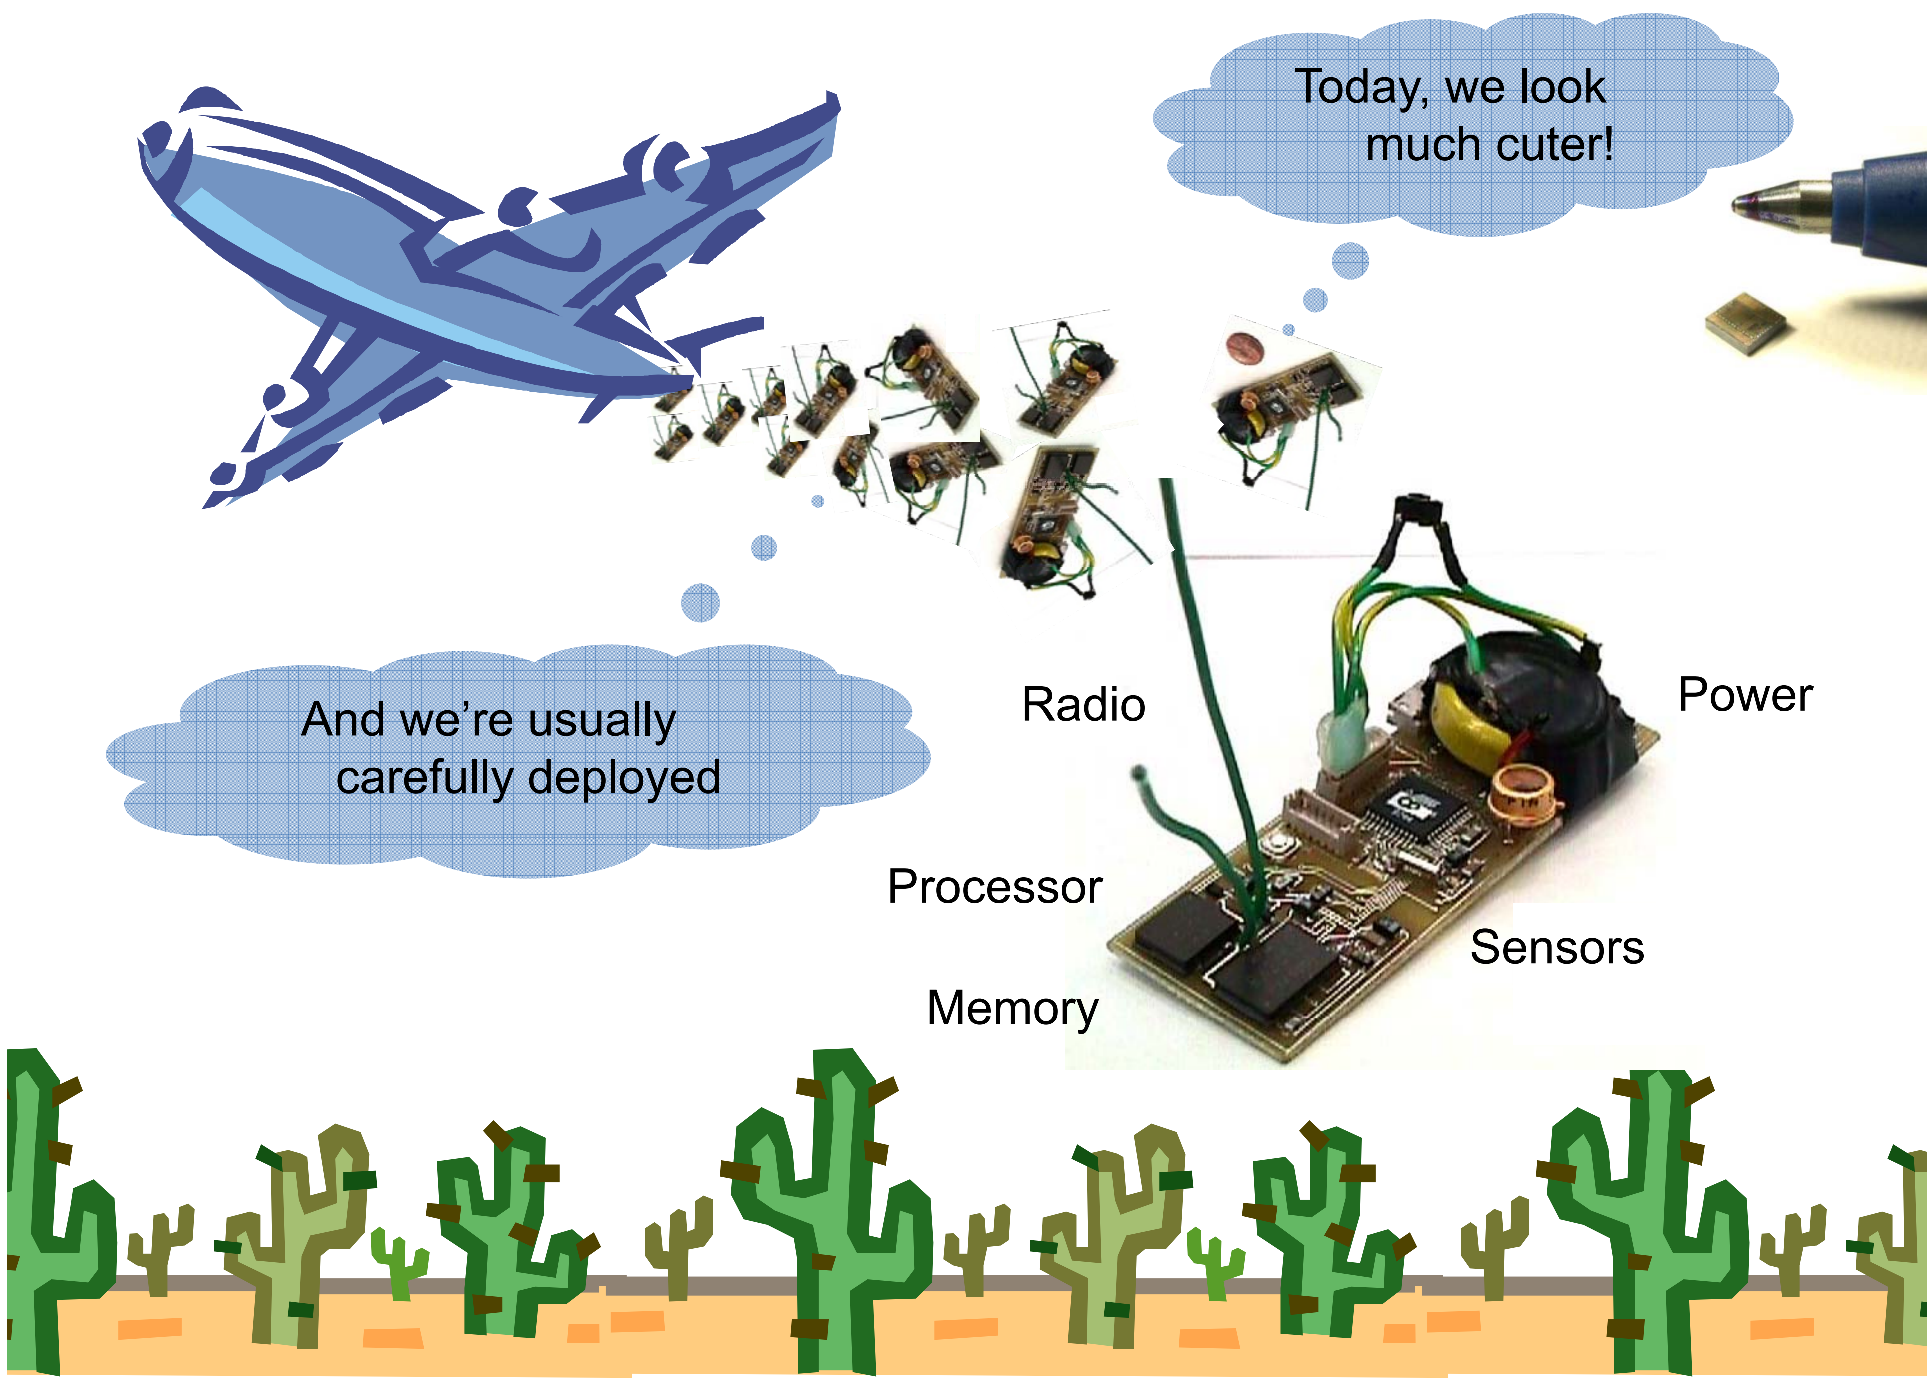
Today, we look much cuter!
And we're usually carefully deployed
Radio
Power
Processor
Sensors
Memory

Periodic data gathering in sensor networks

- All nodes produce relevant information about their vicinity periodically.
- Data is conveyed to an information sink for further processing.
- Data may or may not be aggregated.
- Variations
  - Sense event (e.g. fire, burglar)
  - SQL-like queries (e.g. TinyDB)

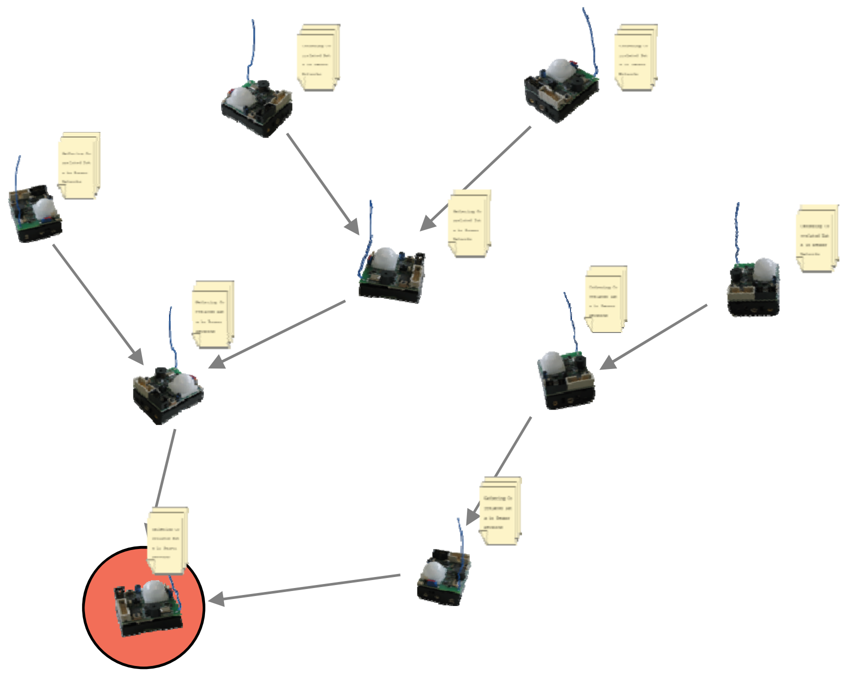

Data Gathering in Wireless Sensor Networks

- 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 peridically wants to compute a function f_n of sensor data
  - At what rate can this function be computed?

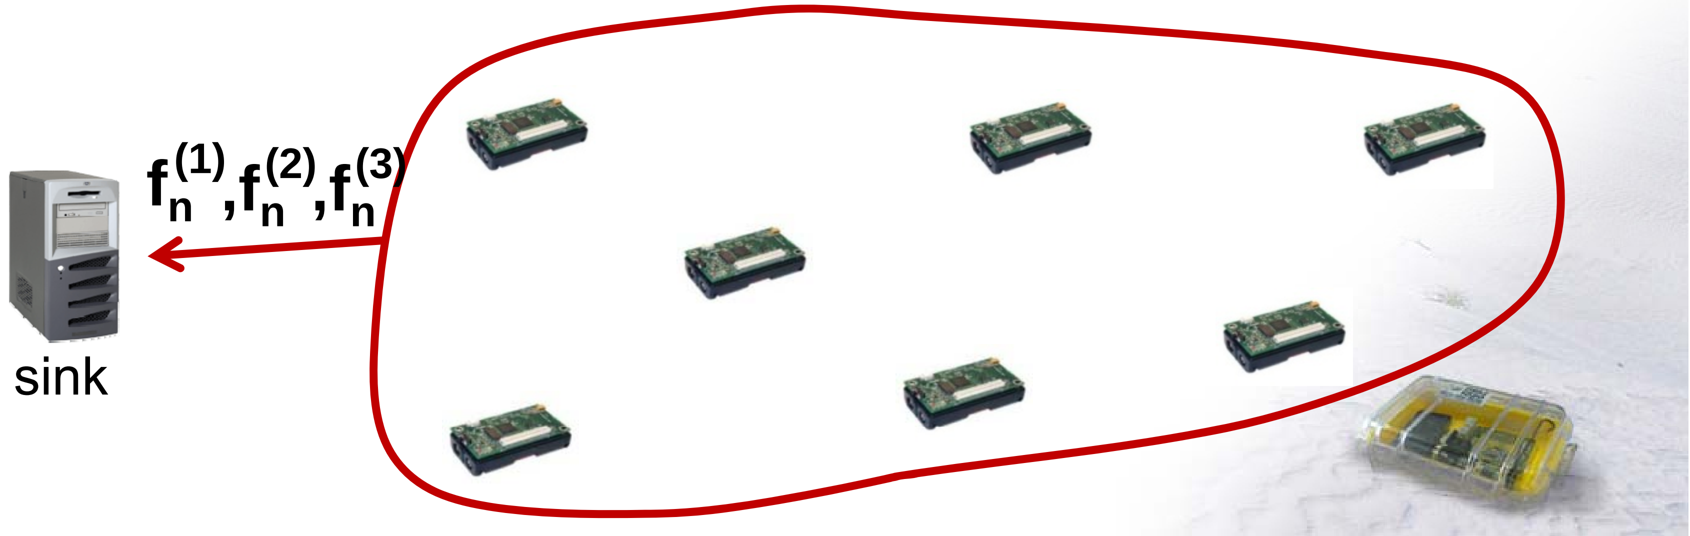

Data Gathering in Wireless Sensor Networks

Example: simple round-robin scheme
→ Each sensor reports its results directly to the root one after another

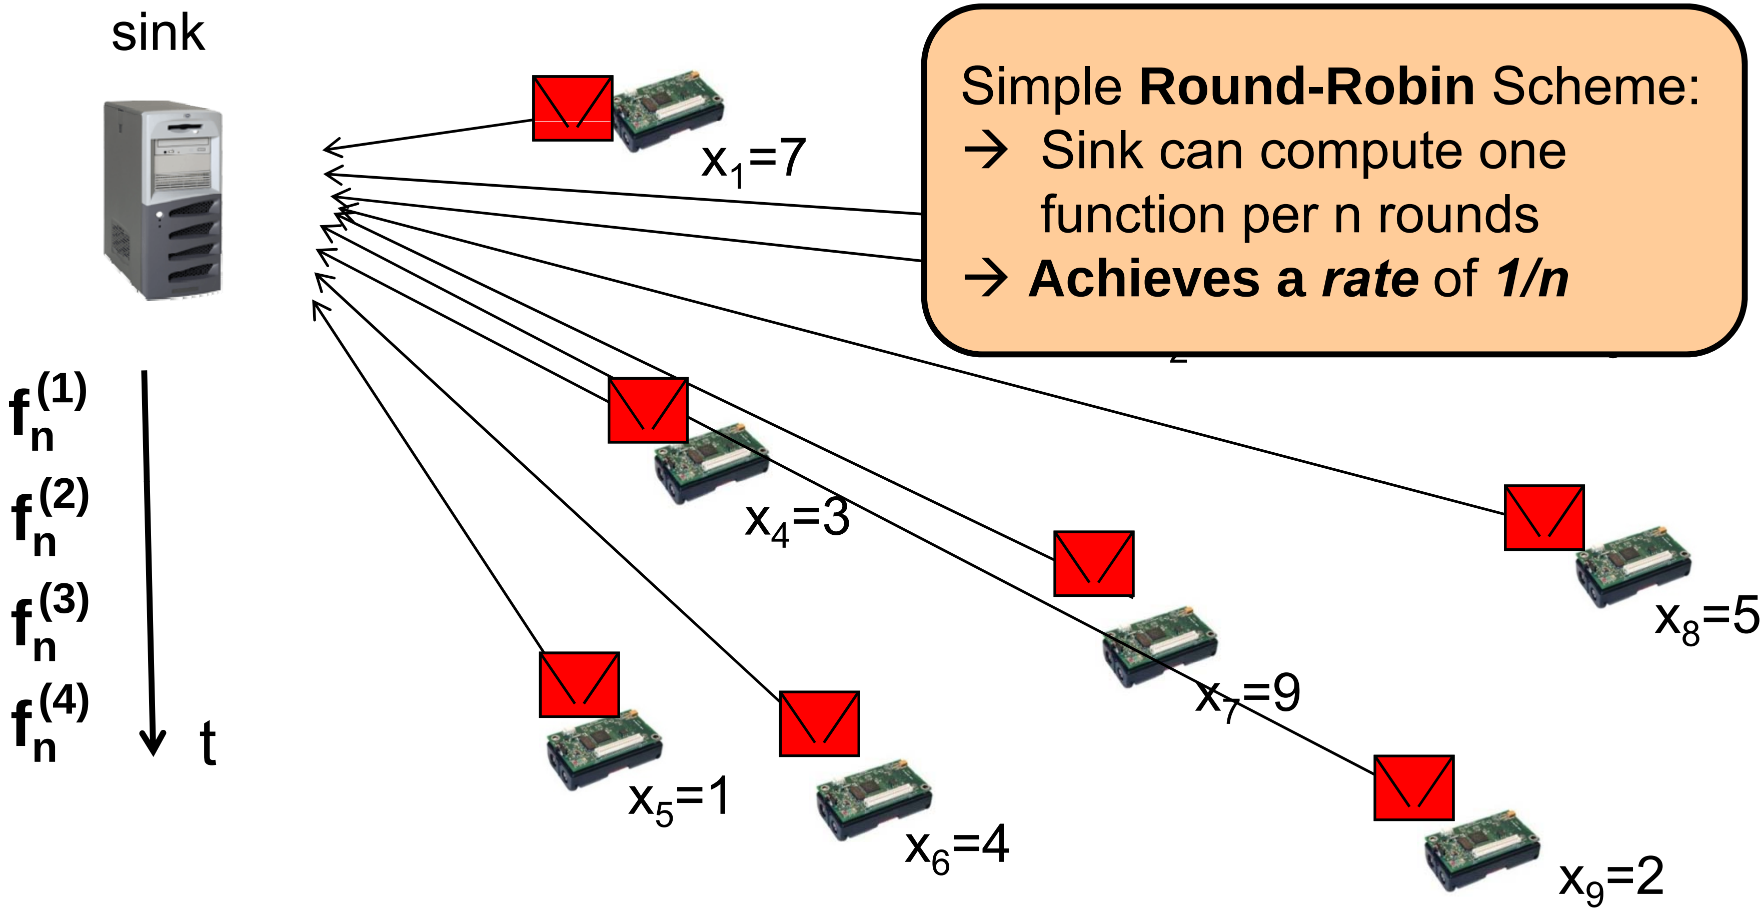

Data Gathering in Wireless Sensor Networks

There are better schemes using
Multi-hop relaying
In-network processing
Spatial Reuse
Pipelining

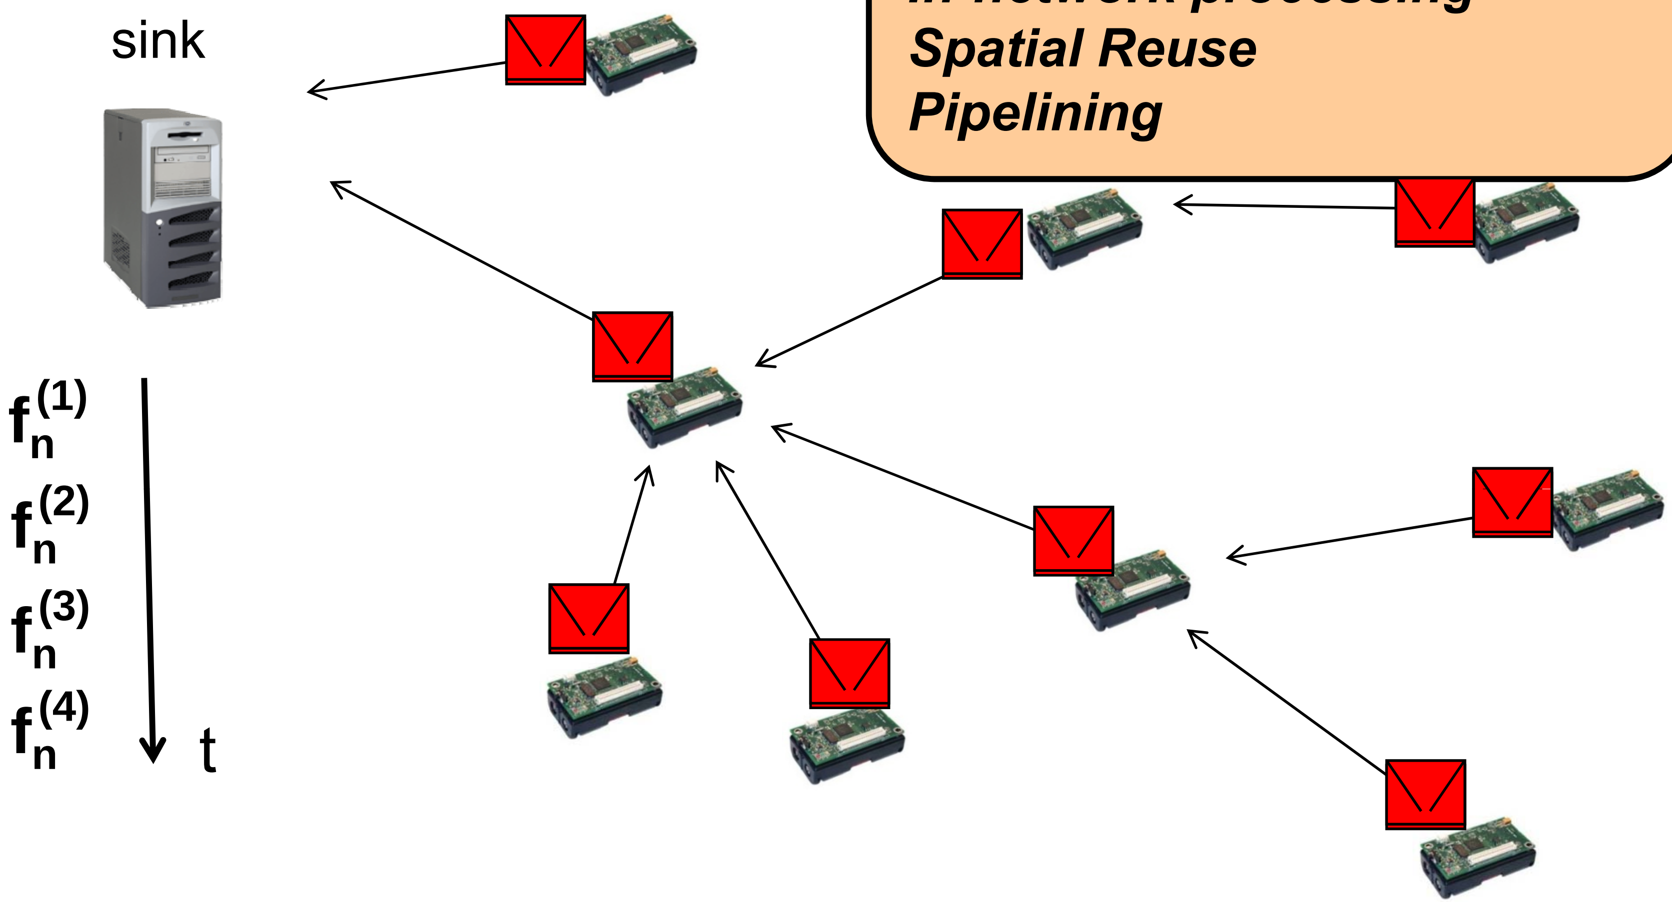

Capacity in Wireless Sensor Networks

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

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

Answer depends on:

1. Function to be computed
2. Coding techniques
3. Network topology
4. Wireless communication model

Function to be computed → Only perfectly compressible functions (max, min, avg,…)

Coding techniques → No fancy coding techniques

"Classic" Capacity…

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]
[Kyasanur 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…

Worst-Case Capacity

- Capacity studies so far make strong assumptions on node deployment, topologies
  - randomly, uniformly distributed nodes
  - nodes placed on a grid
  - etc...

What if a network looks differently…?

Like this?

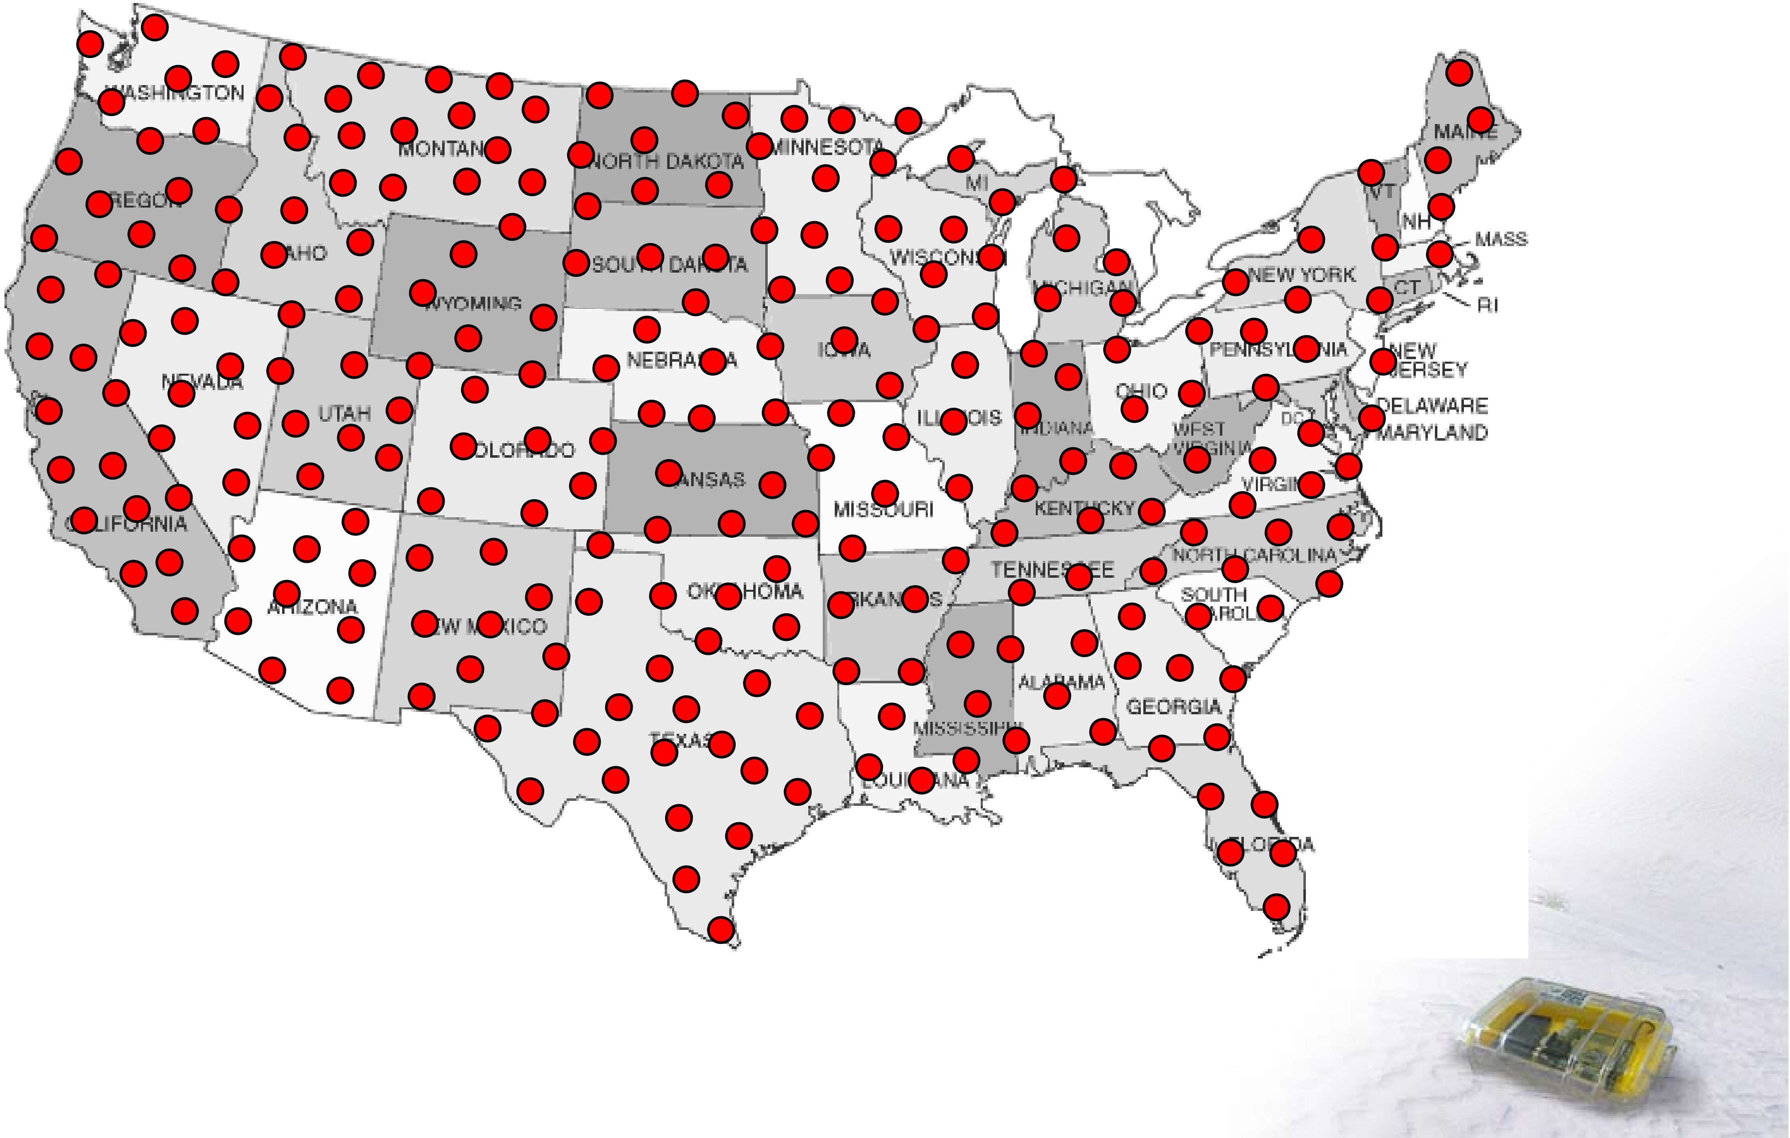
MONTANA
NORTH DAKOTA
MINNESOTA
MI
MAINE
VT
NH
MASS
RI
CT
NEW YORK
WISCONSIN
MICHIGAN
SOUTH DAKOTA
WYOMING
NEBRASKA
IOWA
PENNSYLVANIA
NEW JERSEY
DELAWARE
MARYLAND
NEVADA
UTAH
ILLINOIS
INDIANA
OHIO
WEST VIRGINIA
DC
VIRGINIA
COLORADO
KANSAS
MISSOURI
KENTUCKY
CALIFORNIA
TENNESSEE
NORTH CAROLINA
SOUTH CAROLINA
OKLAHOMA
ARKANSAS
ARIZONA
NEW MEXICO
ALABAMA
GEORGIA
MISSISSIPPI
TEXAS
LOUISIANA
FLORIDA

Or rather like this?

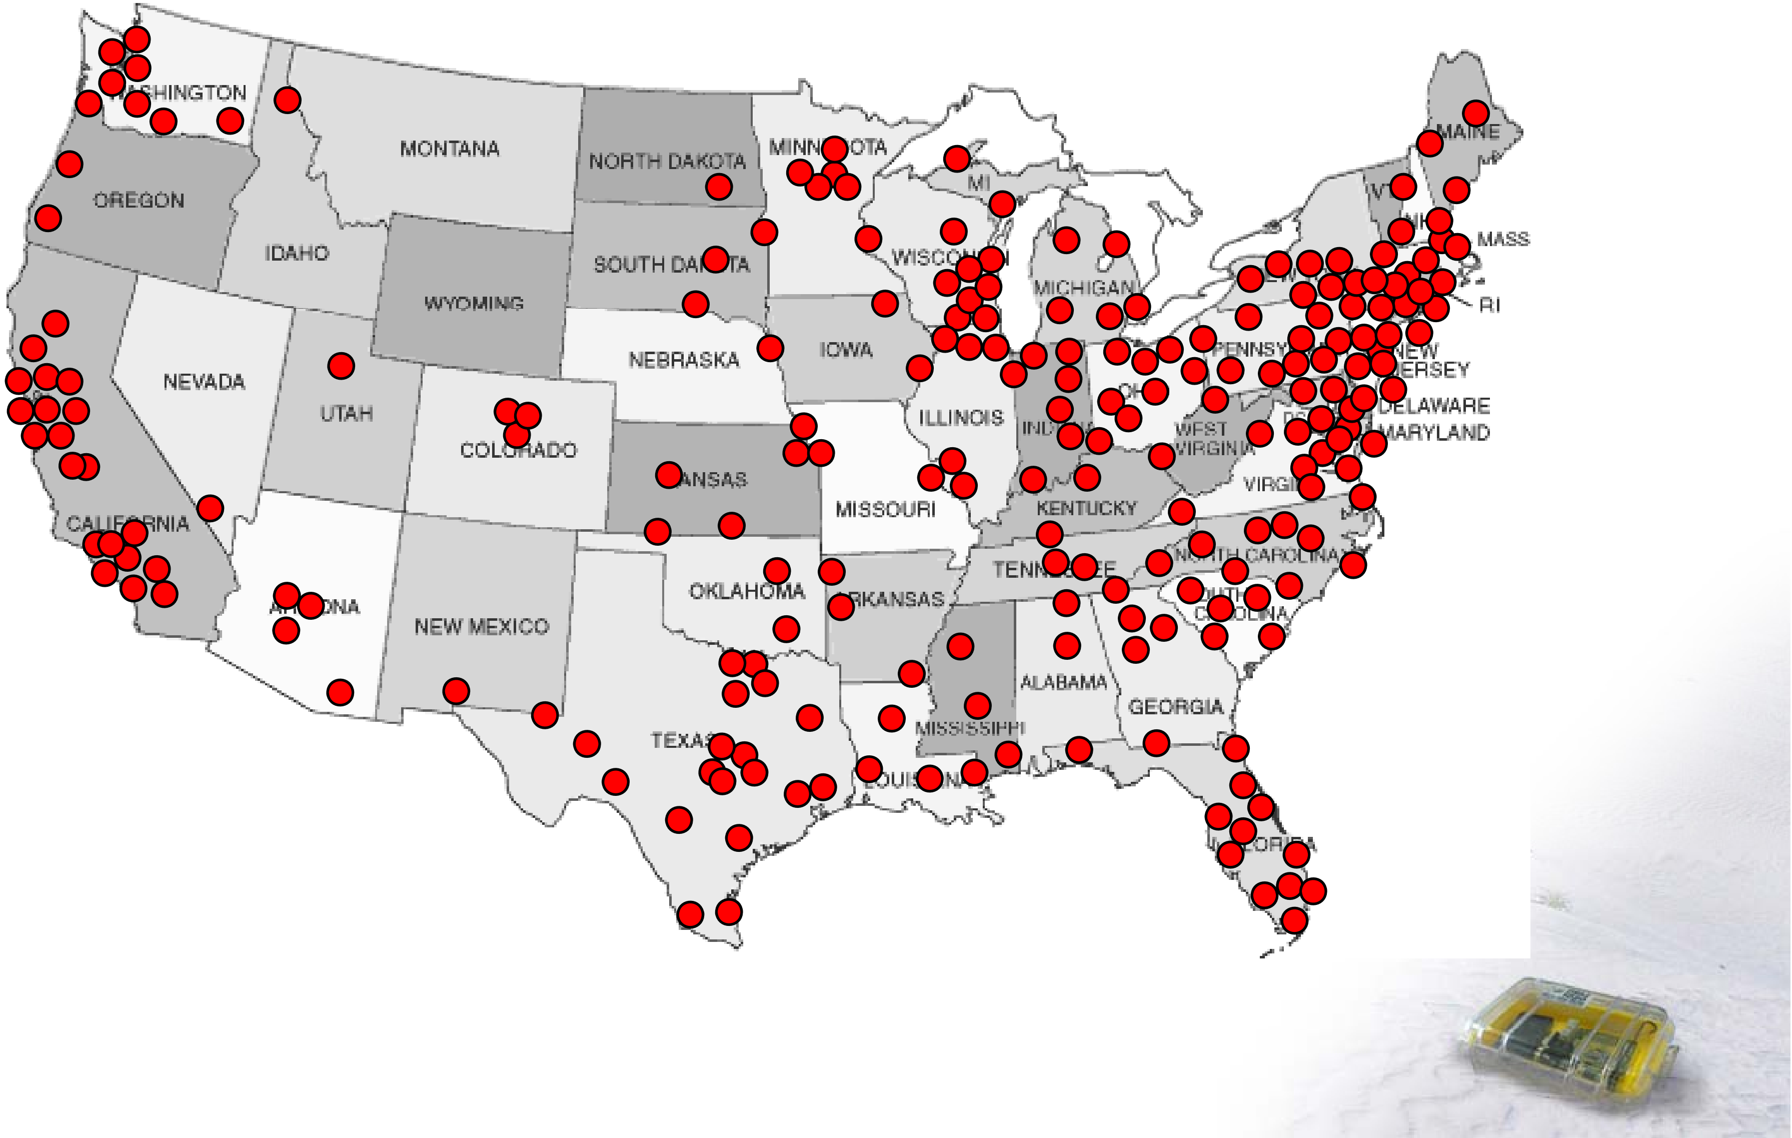

Worst-Case Capacity

- Capacity studies so far have made very strong assumptions on node deployment, topologies
  - randomly, uniformly distributed nodes
  - nodes placed on a grid
  - etc...

What if a network looks differently…?

We assume arbitrary node distribution

worst-case topologies

Classic Capacity	**Worst-Case Capacity**
How much information can be transmitted in nice, well-behaving networks	How much information can be Transmitted in any network

Models

- Two standard models in wireless networking

Protocol Model
(graph-based, simpler)

Physical Model
(SINR-based, more realistic)

Protocol Model

- Based on graph-based notion of interference
- Transmission range and interference range

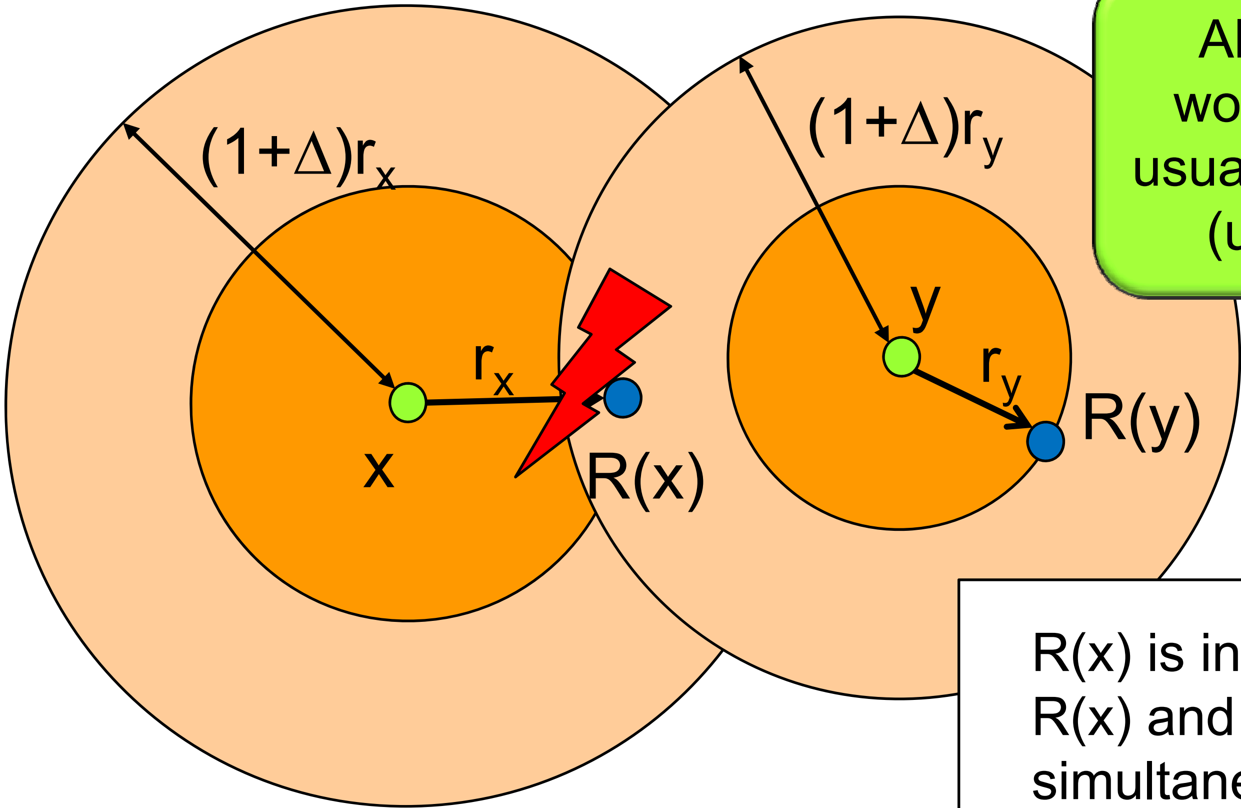

Algorithmic work on worst-case topologies usually in protocol models (unit disk graph,…)

R(x) is in interference range of y
R(x) and R(y) cannot simultaneously receive!

Physical Model

- Based on signal-to-noise-plus-interference (SINR)
- Simplest case:

  → packets can be decoded if SINR is larger than β at receiver

Received signal power from sender

Power level of sender u

Path-loss exponent

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

Minimum signal-to-interference ratio

Noise

Received signal power from all other nodes (=interference)

Distance between two nodes

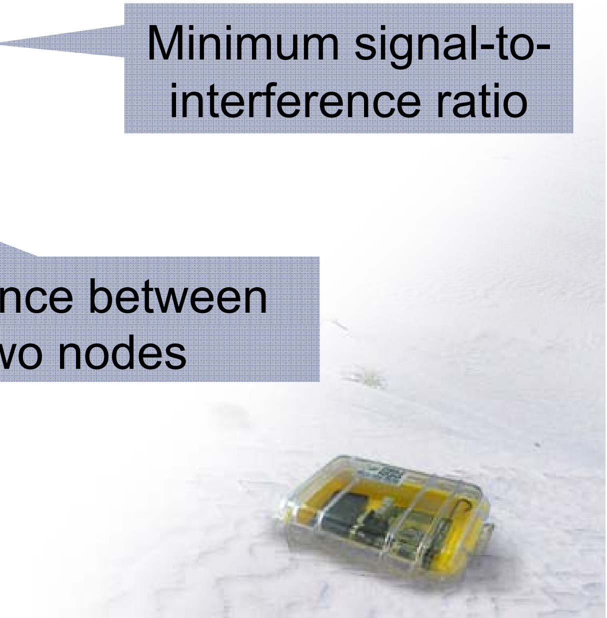

Models

- Two standard models of wireless communication

Protocol Model
(graph-based, simpler)

↔

Physical Model
(SINR-based, more realistic)

- Algorithms typically designed and analyzed in protocol model

Premise: Results obtained in protocol model do not divert too much from more realistic model!

Justification:

Capacity results are typically (almost) the same in both models
(e.g., Gupta, Kumar, etc...)

Example: Protocol vs. Physical Model

A sends to D, B sends to C

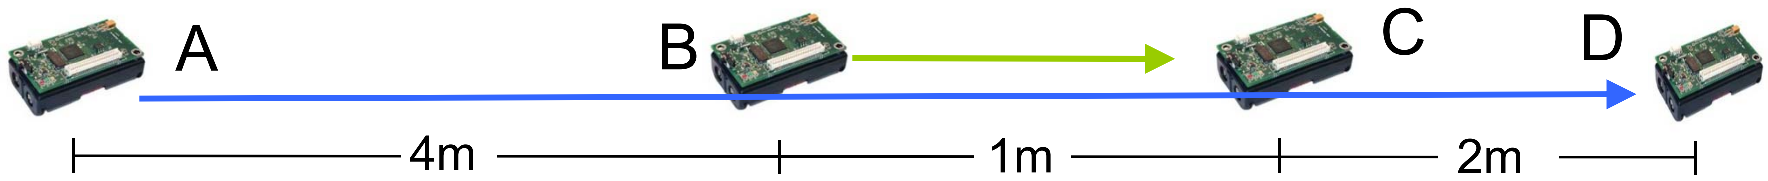

Assume a single frequency (and no fancy decoding techniques!)

Is spatial reuse possible? → NO — Protocol Model

Is spatial reuse possible? → YES — Physical Model

YES — In Reality!

Let α=3, β=3, and N=10nW

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

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

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

This works in practice!

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

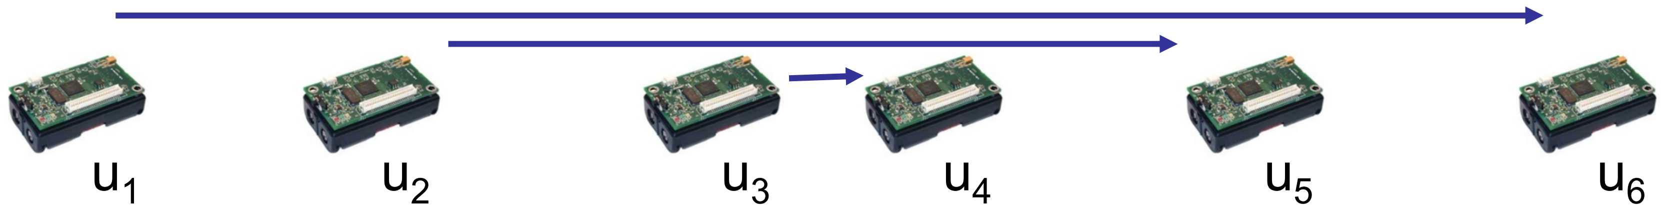

- Time for successfully transmitting 20‘000 packets:

	Time required	
	standard MAC	“SINR-MAC”
Node u_1	$721s$	$267s$
Node u_2	$778s$	$268s$
Node u_3	$780s$	$270s$

	Messages received	
	standard MAC	“SINR-MAC”
Node u_4	19999	19773
Node u_5	18784	18488
Node u_6	16519	19498

Speed-up is almost a factor 3

[Moscibroda, Wattenhofer, Weber, Hotnets‘06]

Upper Bound Protocol Model

- There are networks, in which at most one node can transmit!
  → like round-robin
- Consider exponential node chain
- Assume nodes can choose arbitrary transmission power

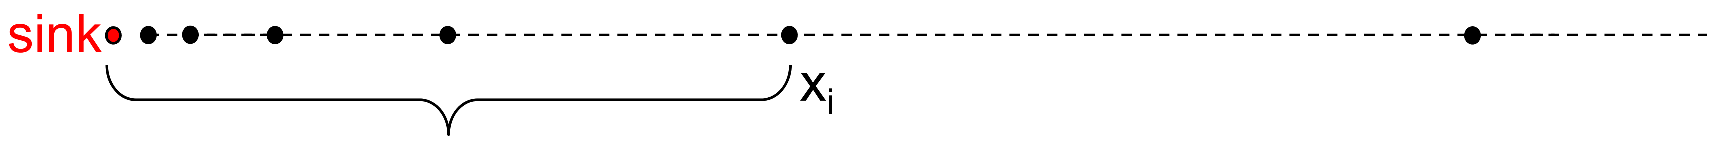

$d(\text{sink},x_i) = (1+1/\Delta)^{i-1}$

- Whenever a node transmits to another node
  → All nodes to its left are in its interference range!
  → Network behaves like a single-hop network

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

Lower Bound Physical Model

- Much better bounds in SINR-based physical model are possible (exponential gap)
- Paper presents a scheduling algorithm that achieves a rate of $\Omega(1/\log^3 n)$

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

- Algorithm is centralized, highly complex → 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 → yields forest
4) Only root of each tree remains active

loop until no active nodes

Can be adjusted if transmission power limited

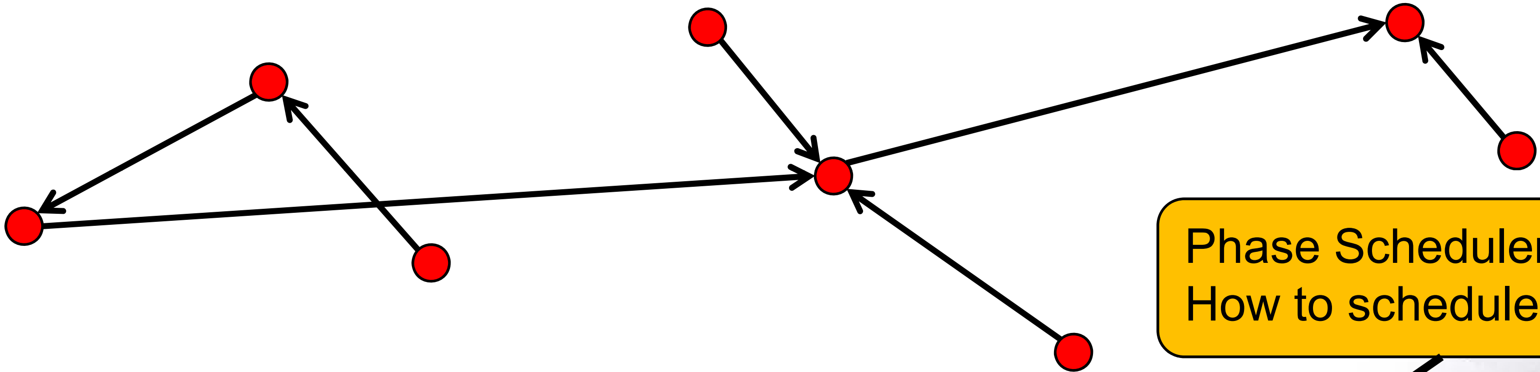

Phase Scheduler:
How to schedule T(X)?

The resulting structure has some nice properties
→ If each link of T(X) can be scheduled at least once in L(X) time-slots
→ 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!

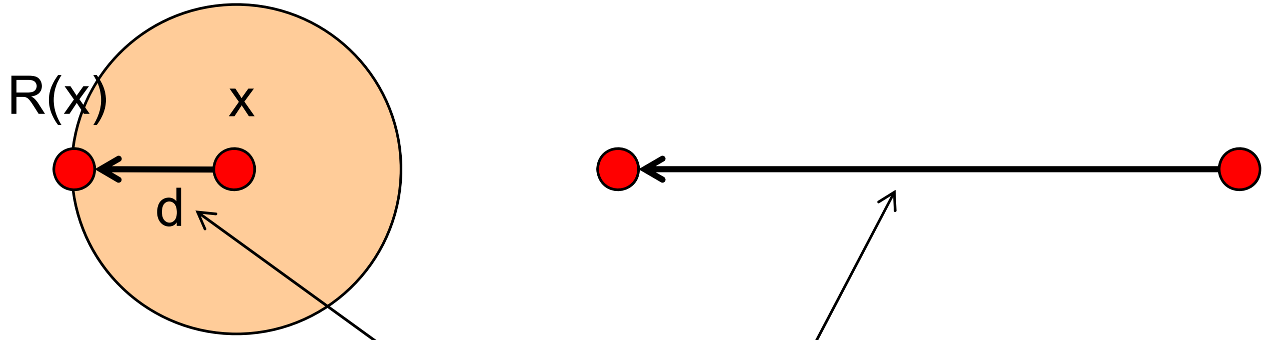

If we want to schedule both links…

1) … R(x) must be overpowered
→ Must transmit at power more than ~d^{α}

2) If senders of small links overpower their receiver…
… their "safety radius" increases (spatial reuse smaller)

Subtle balance is needed!

Scheduling Algorithm – Phase Scheduler

1) Partition links into sets of similar length

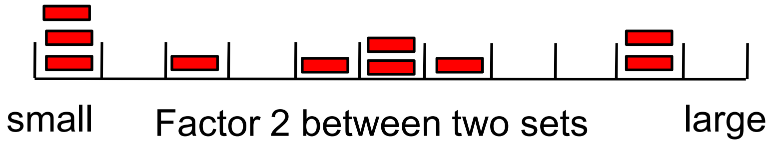

2) Group sets such that links a and b in two sets in the same group have at least $d_a \geq (\xi\beta)^{\xi(\tau a-\tau b)} \cdot d_b$

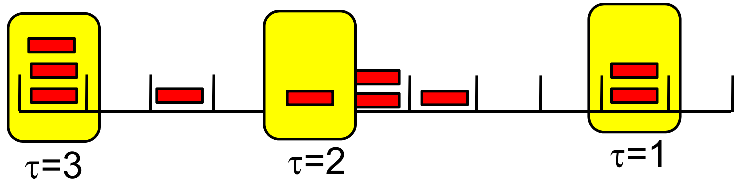

→ Each link gets a τ_{ij} value → Small links have large τ_{ij} and vice versa
→ Schedule links in these sets in one outer-loop iteration
→ Intuition: Schedule links of similar length or very different length

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) → $\Omega(1/\log^3 n)$ bound

Worst-Case Capacity in Wireless Networks

Worst-Case Capacity

Traditional Capacity

Model \ Networks	Max. rate in arbitrary, worst-case deployment	Max. rate in random, uniform deployment
Protocol Model	$\Theta(1/n)$	$\Theta(1/\log n)$
Physical Model	$\Omega(1/\log^3 n)$	$\Omega(1/\log n)$

[Giridhar, Kumar, 2005]

Exponential gap between protocol and physical model!

<u>The Price of Worst-Case Node Placement</u>

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

Possible Applications – Improved "Channel Capacity"

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

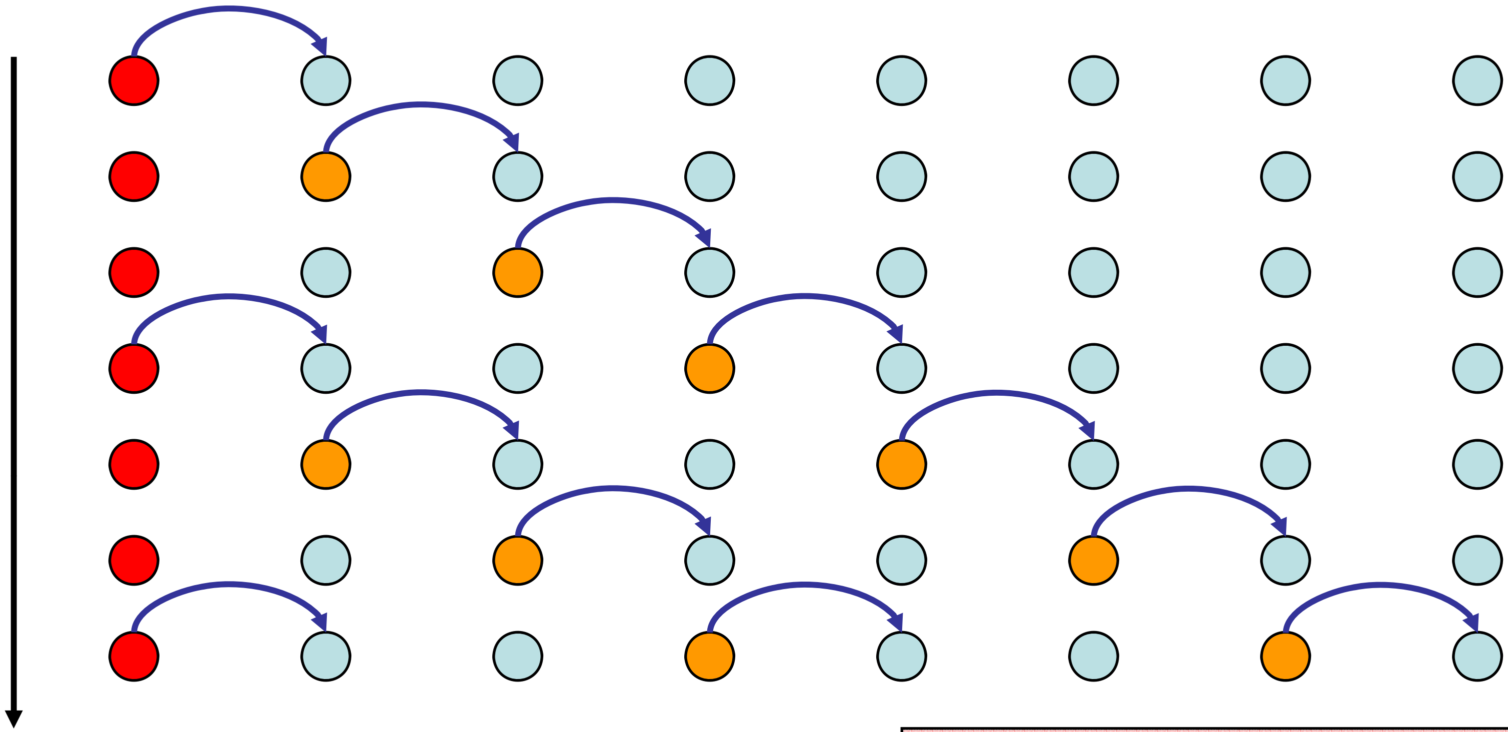

time

Channel capacity is 1/3

Possible Applications – Improved “Channel Capacity”

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

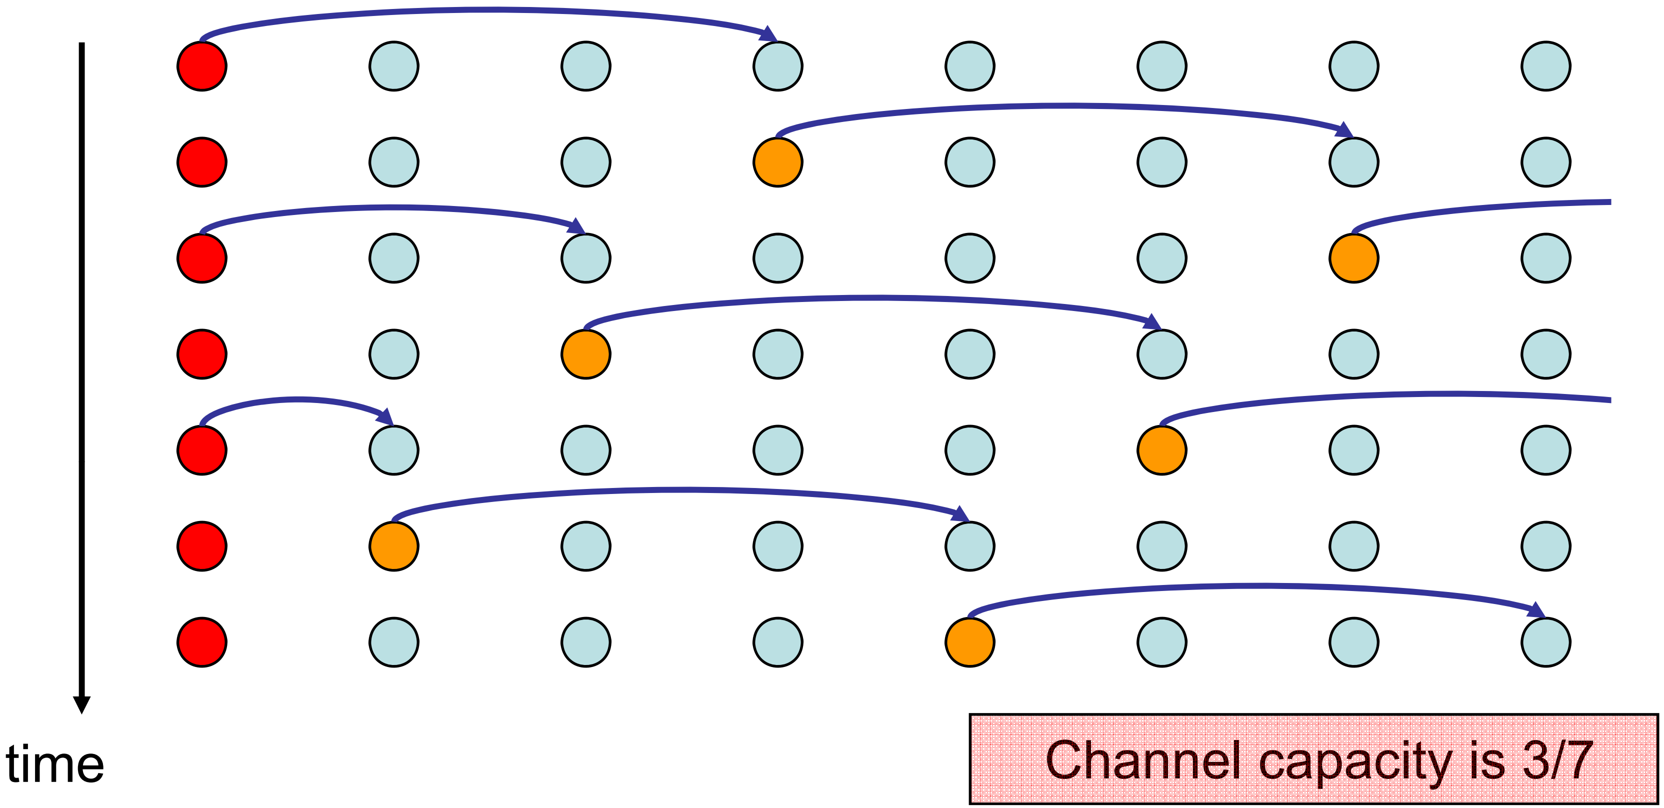

Possible Applications – Improved "Channel Capacity"

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

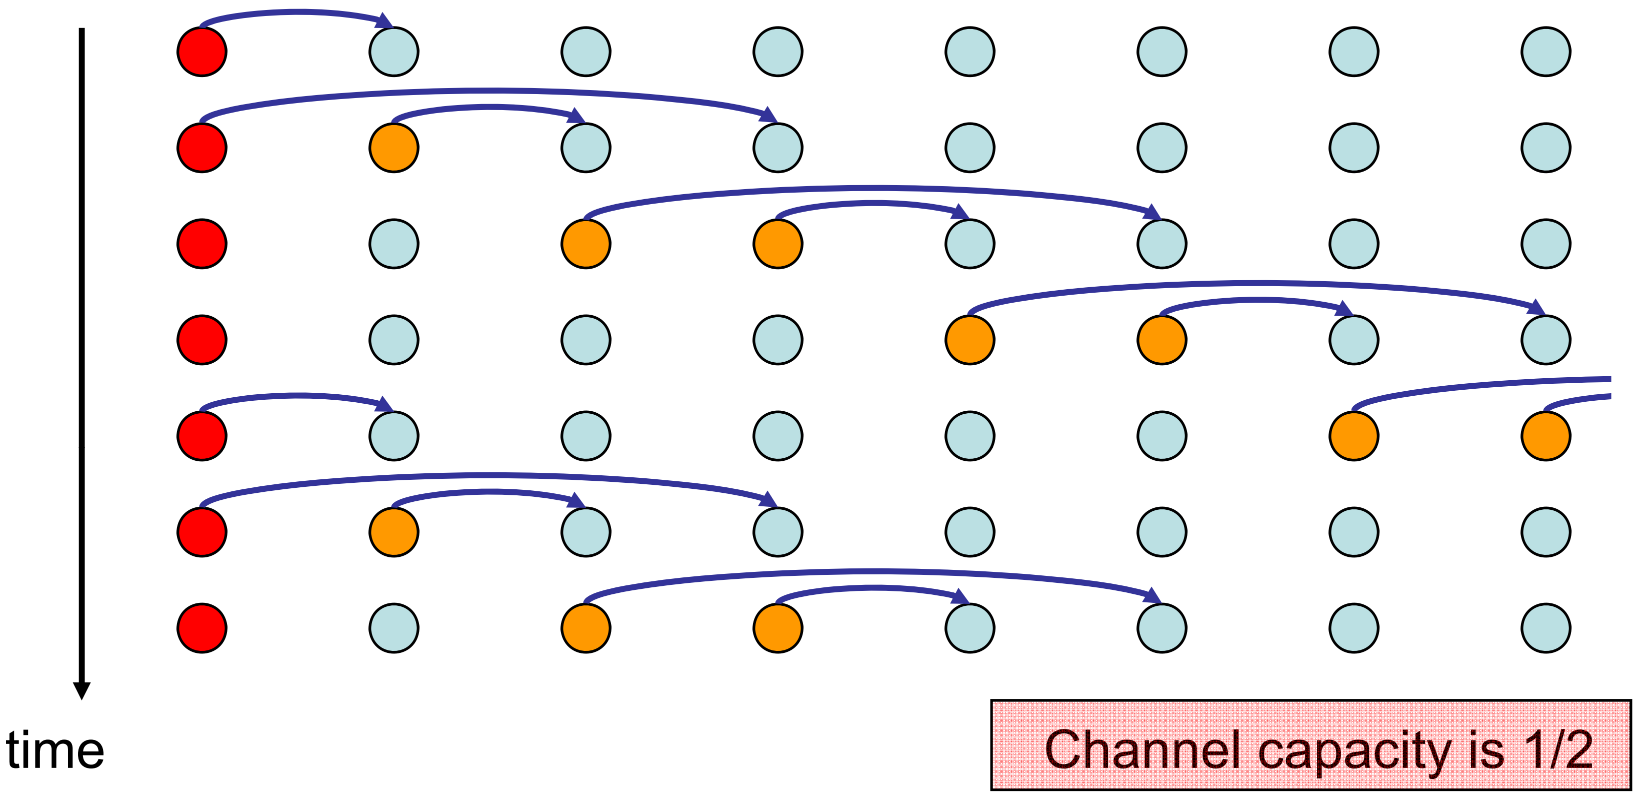

Possible Application – Hotspots in WLAN

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

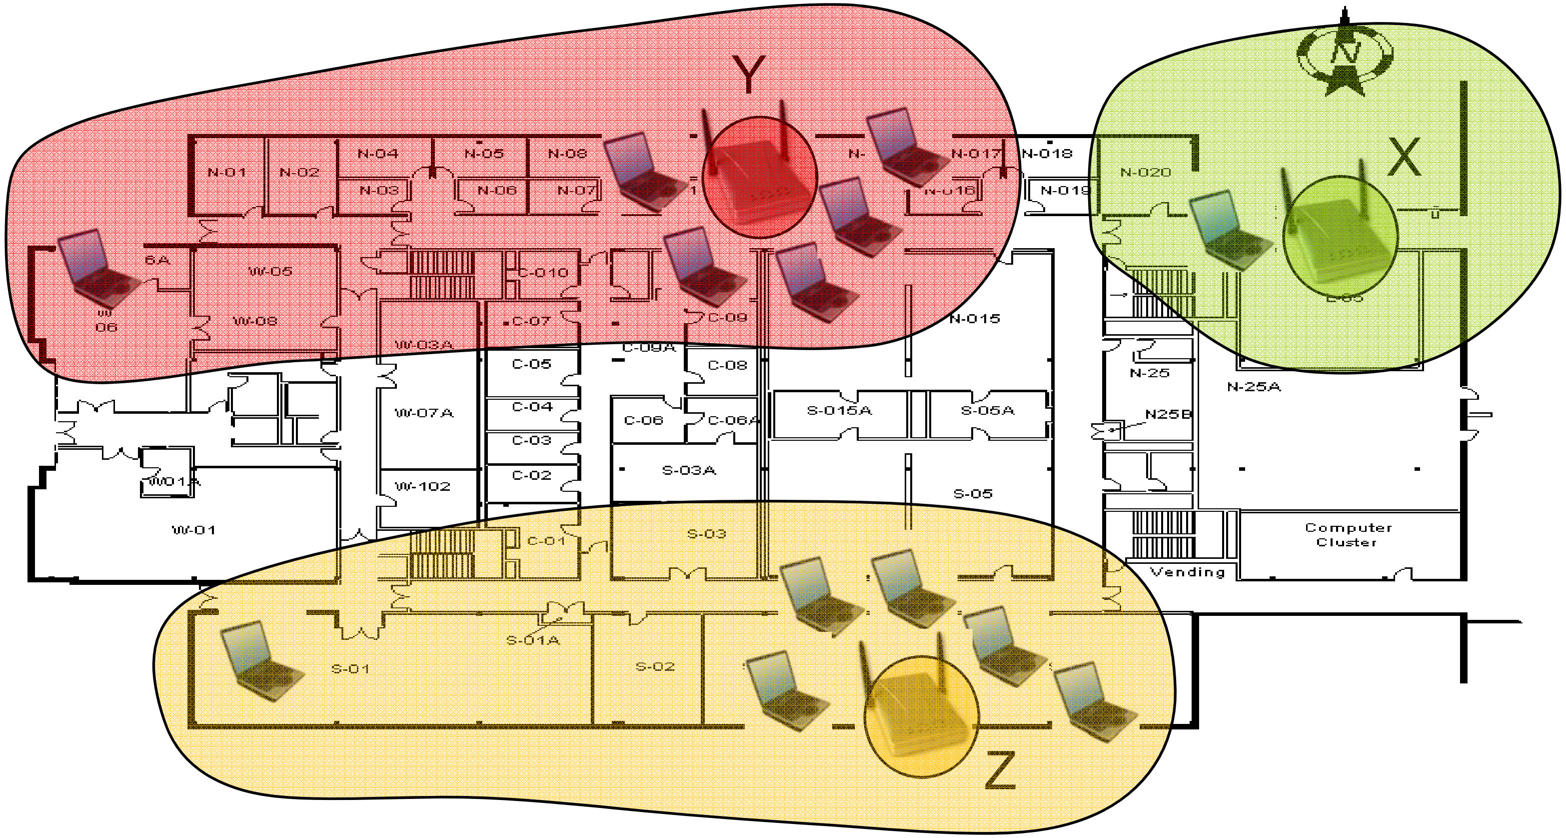

Possible Application – Hotspots in WLAN

- 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!

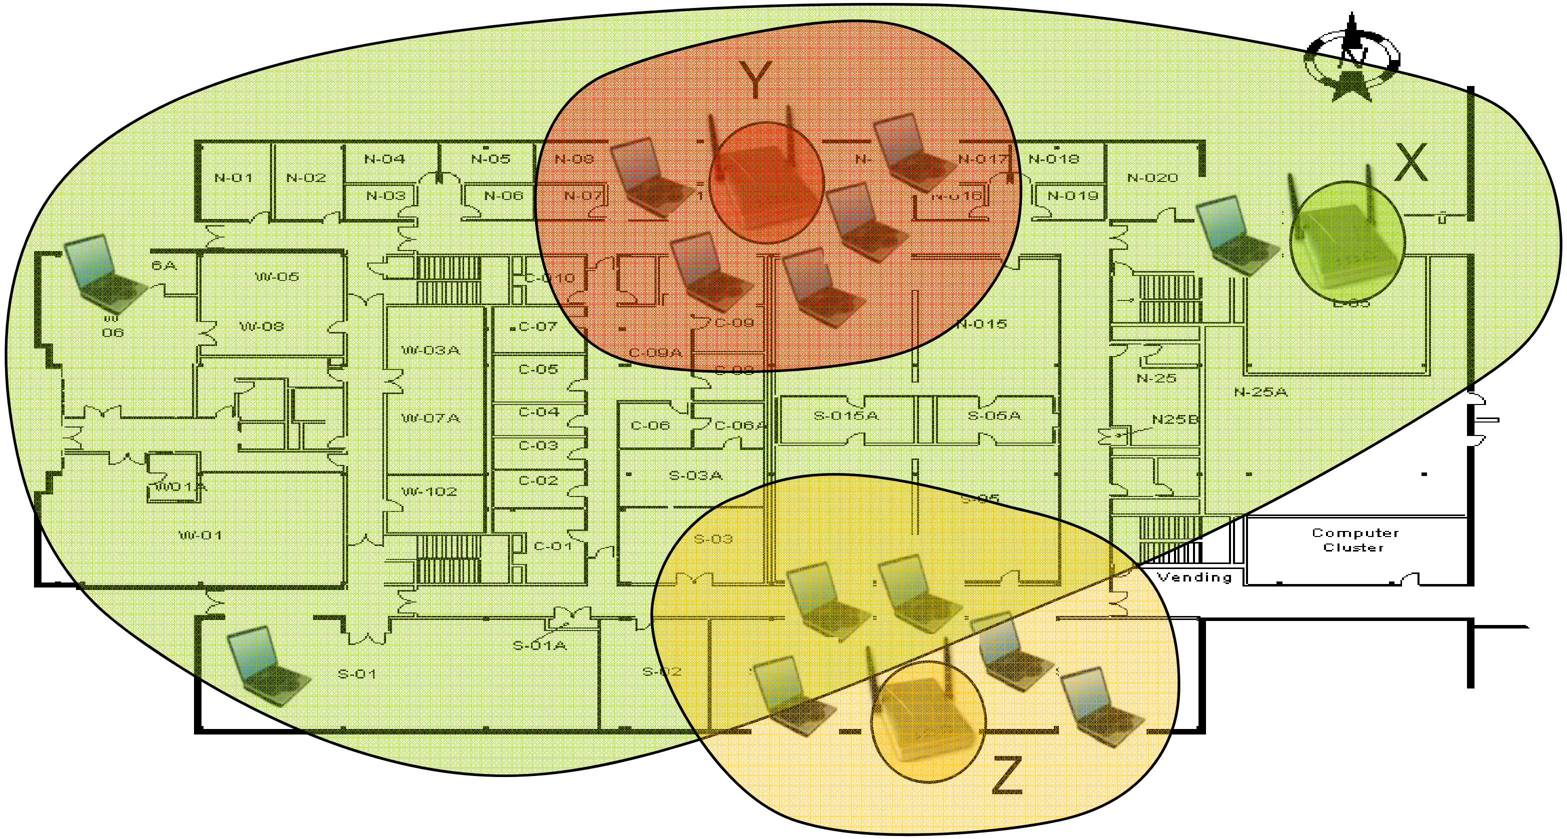

Possible Applications – Data Gathering

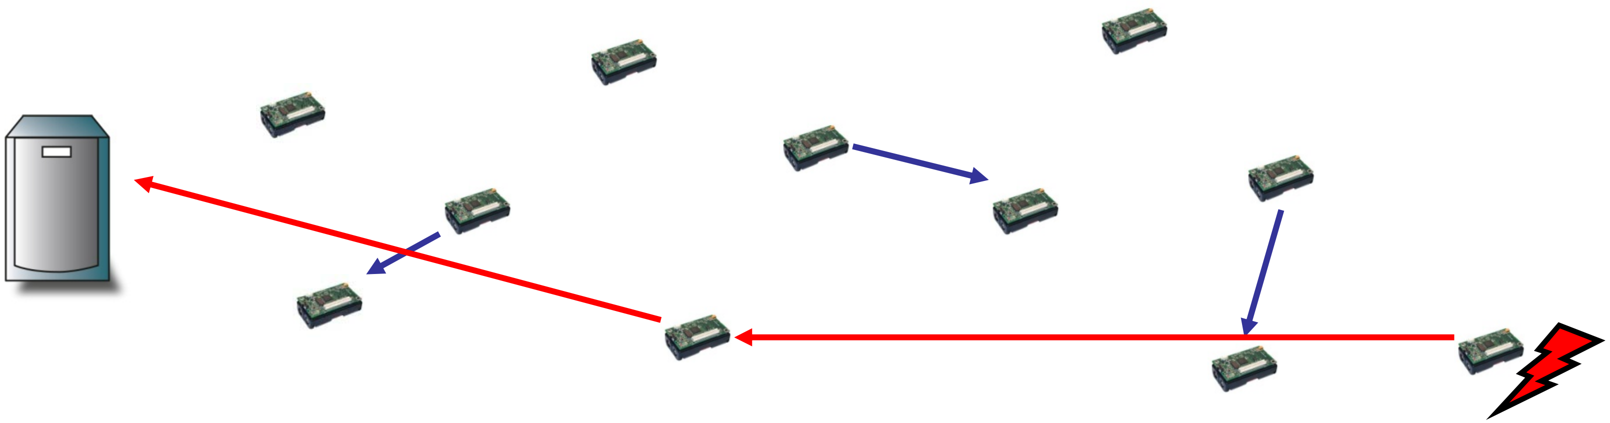

- Neighboring nodes must communicate periodically
  (for time synchronisation, neighborhood detection, etc…)
- Sending data to base station may be time critical → use long links
- Employing clever power control may reduce delay & reduce coordination overhead!

→ From theory (scheduling) to practice (protocol design)…?

Summary

- Introduce worst-case capacity of sensor networks
  → How much data can periodically be sent to data sink
- Complements existing capacity studies
- Many novel insights

1) Possibilities and limitations of wireless communication
2) Fundamentals of wireless communication models
3) How to devise efficient scheduling algorithms, protocols

Sensor Networks Scale!
Efficient data gathering is possible in every (even worst-case) network!

Protocol Model Poor!
Exponential gap between protocol and physical model!

Efficient Protocols!
Must use SINR-effects and power control to achieve high rate!

Remaining Questions…?

- My talk so far was based on the paper Moscibroda & W, The Complexity of Connectivity in Wireless Networks, Infocom 2006

- The paper was more general than my presentation
  - It was not about data gathering rate, but rather…
  1. Given an arbitrary network
  2. Connect the nodes in a meaningful way by links
  3. Schedule the links such that the network becomes strongly connected

- Question: Given *n* communication requests, assign a color (time slot) to each request, such that all requests sharing the same color can be handled correctly, i.e., the SINR condition is met at all destinations (the source powers areconstant). The goal is to minimize the number of colors.

Is this a difficult problem?

Scheduling Wireless Links: How hard is it?

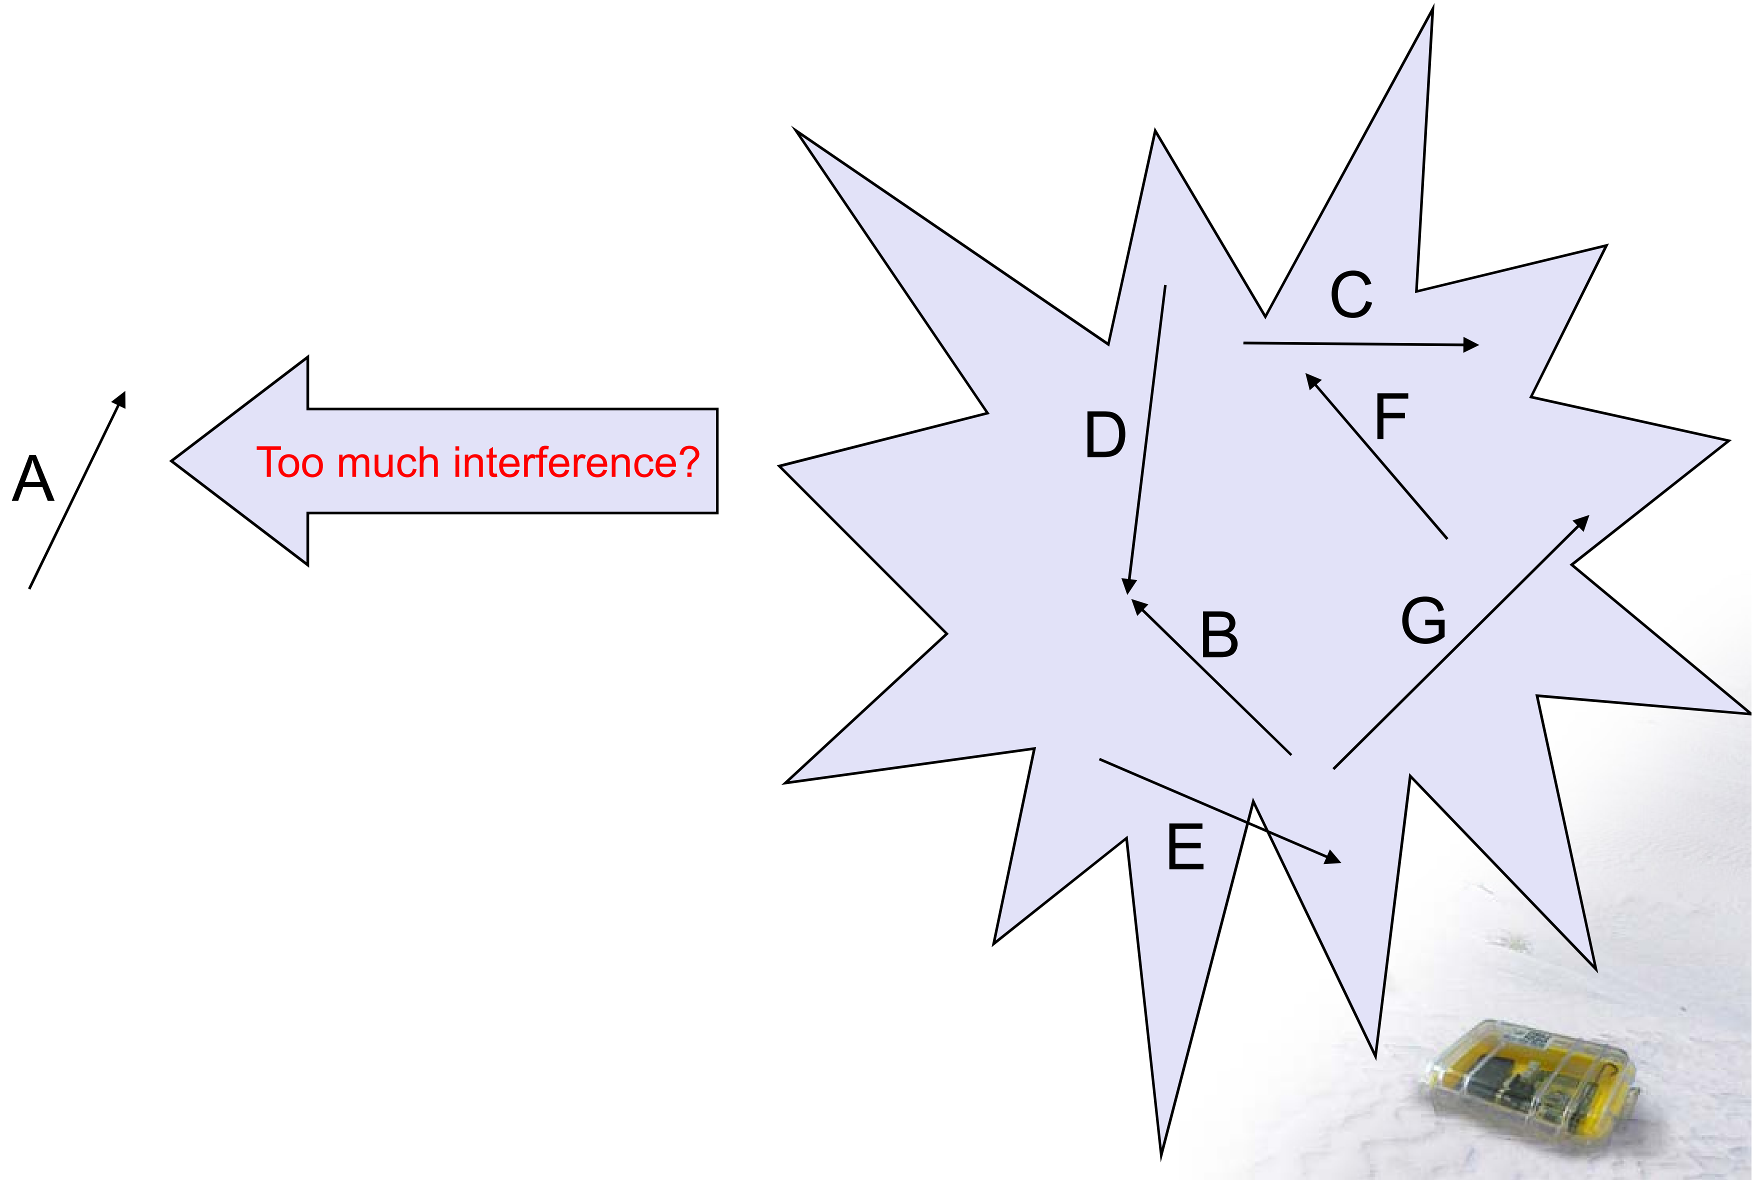

Scheduling: Problem Definition

- *P*: constant power level
- *L*: set of communication requests
- *S*: schedule $S = \{S_1, S_2, \ldots, S_T\}$

- Interference Model: *SINR*
  - *A*: path-loss matrix, defined for every pair of nodes

Received signal power from sender

Min. SINR threshold

$$SINR(s,r) = \frac{\frac{P}{A_{sr}}}{N + \sum_{v \in V, v \neq s} \frac{P}{A_{vr}}} \geq \beta$$

Ambient noise

Received signal power from all other nodes (Interference!)

- **Problem statement:**

> *Find a minimum-length schedule **S**, s.t. every link in **L** is scheduled in at least one time slot **t, 1≤t ≤T**, and all concurrently scheduled receivers in $\boldsymbol{S_t}$ satisfy the **SINR** constraints.*

“Scheduling as hard as coloring” … not really!

“The Wall Model”: Now only adjacent links interfere! (Has been shown to be as hard as coloring [Bjoerklund 2003])

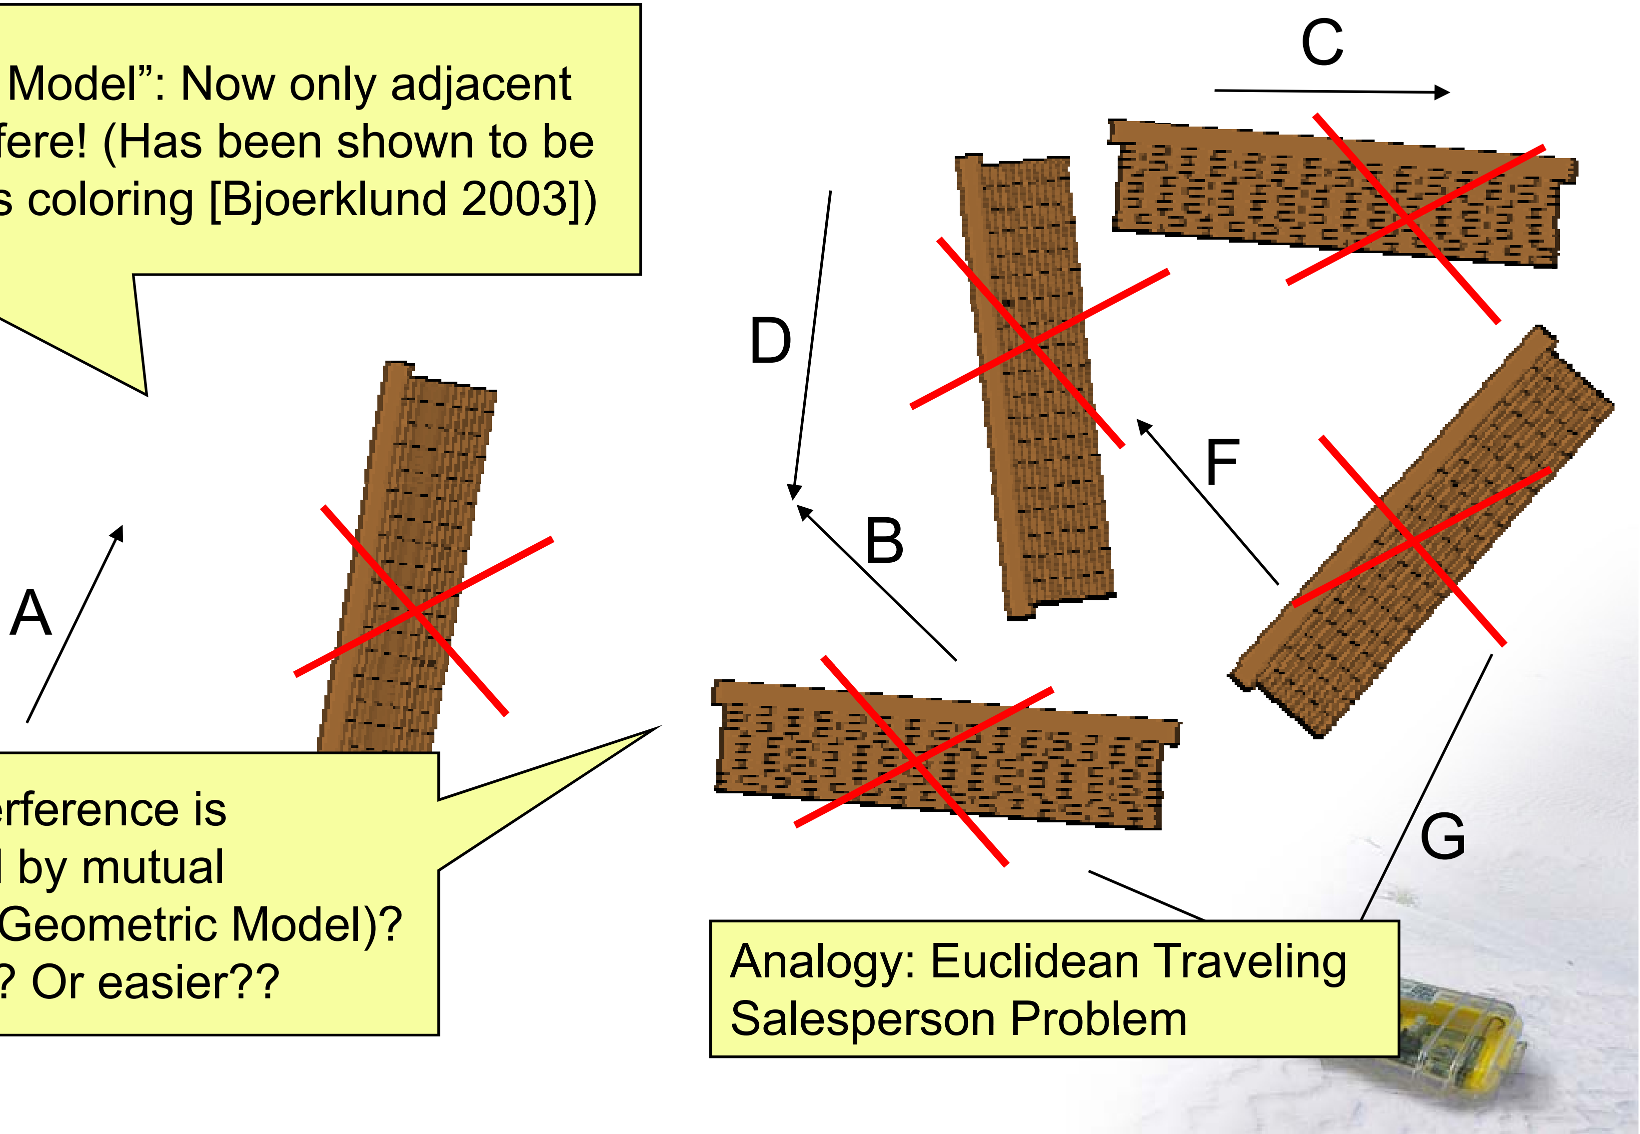

What if interference is determined by mutual distances (Geometric Model)? Is it harder? Or easier??

Analogy: Euclidean Traveling Salesperson Problem

Scheduling: Reduction from Partition

- Partition problem (NP-Complete [Karp 1972]):
  - Given a set of integers I, find two subsets of integers I_1, I_2, s.t.:
- Decision version of Scheduling: $T \le 2$:
  - Consider a set of integers I, whose elements sum up to σ:

$$I_1, I_2 \subset I = \{i_1, \ldots, i_n\}$$
$$I_1 \cap I_2 = \emptyset,$$
$$I_1 \cup I_2 = I,$$
$$\sum_{i_j \in I_1} i_j = \sum_{i_j \in I_2} i_j = \frac{1}{2} \sum_{i_j \in I} i_j.$$

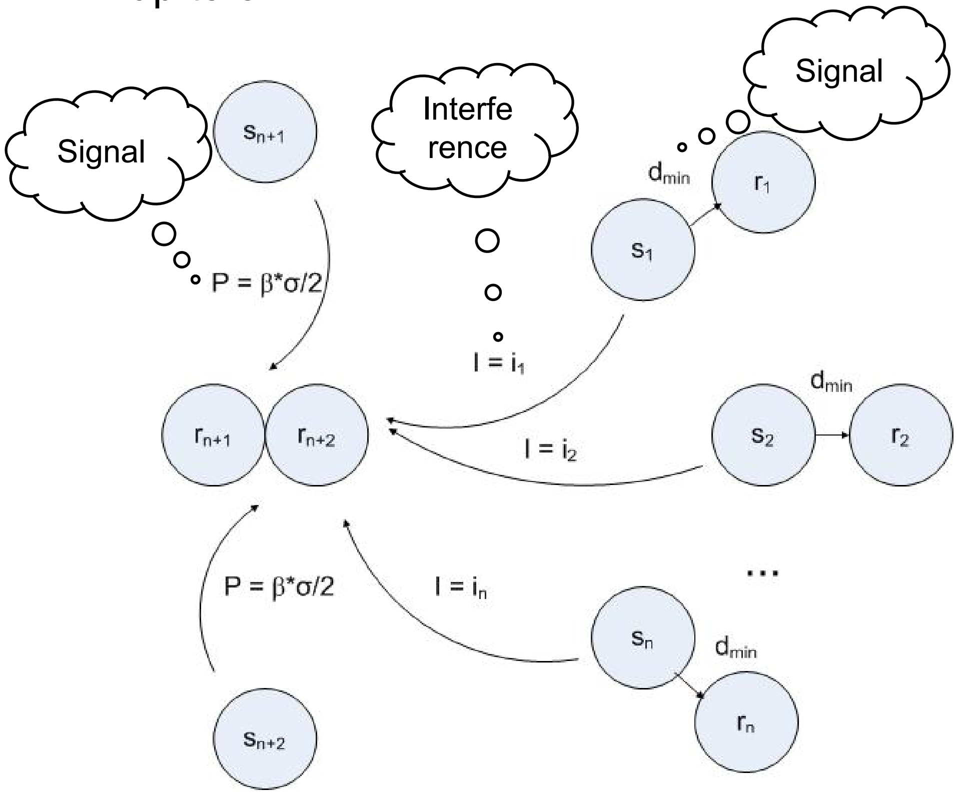

$$SINR_{r_{n+1}} = \frac{\beta \cdot \frac{\sigma}{2}}{\sum_{i_j \in I_1} i_j} = \beta$$

Schedule with time $T \le 2 \leftrightarrow$ Partition

SINR Models

- Abstract SINR
  - Arbitrary path loss matrix
  - No notion of triangle inequality
  - If an algorithm works here, it works everywhere!
  - Best model for upper bounds
- Geometric SINR
  - Nodes are points in plane
  - Path loss is function of distance
  - If an impossibility result holds here, it holds everywhere!
  - Best model for lower bounds

too pessimistic **too optimistic**

- Reality is here
  - Path loss roughly follows geometric constraints, but there are exceptions
  - Open field networks are closer to Geometric SINR
  - With more walls, you get more and more Abstract SINR

Models can be put in relation

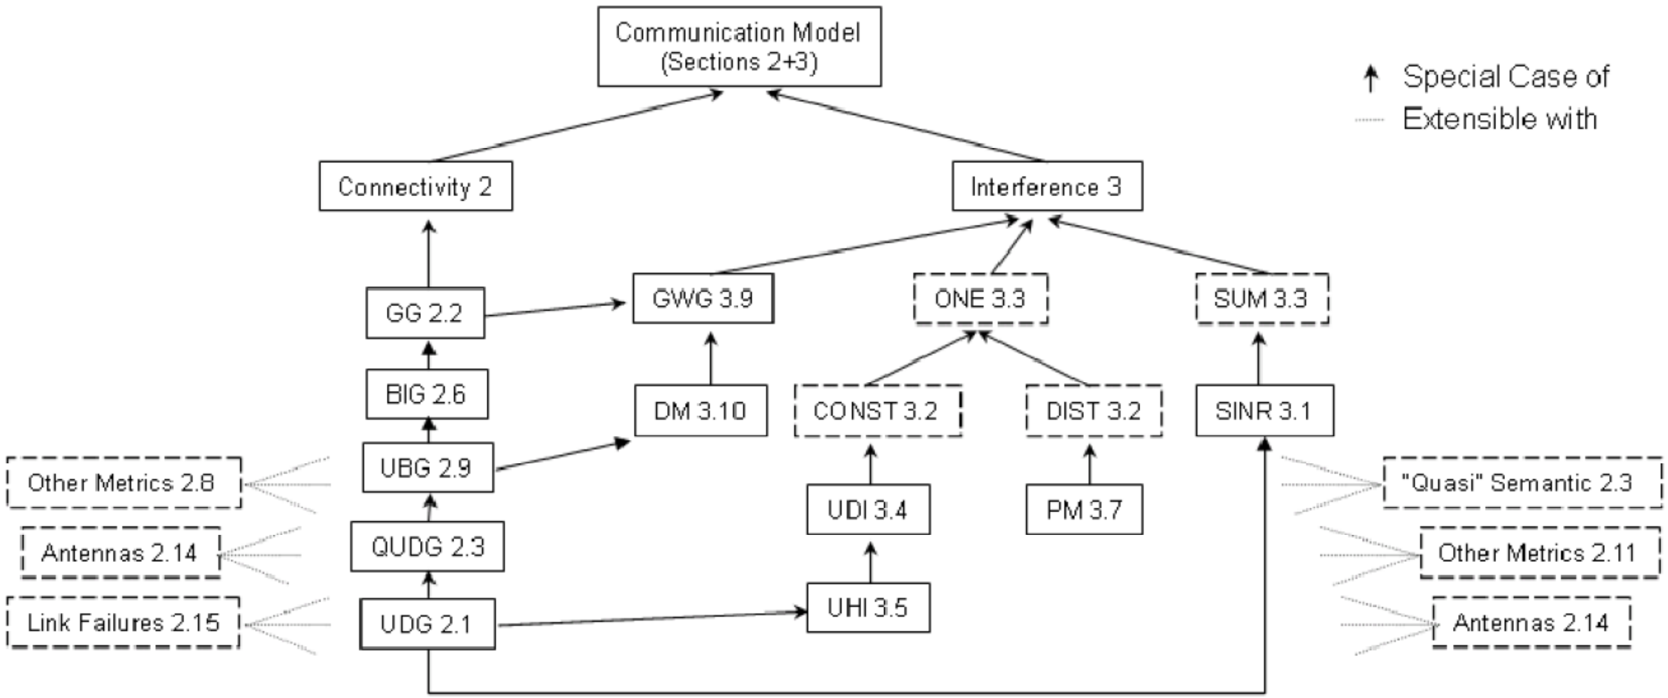

- Try to proof correctness in an as “high” as possible model
- For efficiency, a more optimistic (“lower”) model might be fine
- Lower bounds are best proved in “low” models

Overview of results so far

- Moscibroda, W, Infocom 2006
  - First paper in this area, O(log^3 *n*) bound for connectivity, and more
- Moscibroda, W, Weber, HotNets 2006
  - Practical experiments, ideas for capacity-improving protocol
- Goussevskaia, Oswald, W, MobiHoc 2007
  - Hardness results & constant approximation for constant power

- Moscibroda, W, Zollinger, MobiHoc 2006
  - First results beyond connectivity, namely in the topology control domain
- Moscibroda, Oswald, W, Infocom 2007
  - Generalizion of Infocom 2006, proof that known algorithms perform poorly
- Chafekar, Kumar, Marathe, Parthasarathy, Srinivasan, MobiHoc 2007
  - Cross layer analysis for scheduling and routing
- Moscibroda, IPSN 2007
  - Connection to data gathering, improved O(log^2 *n*) result
- Goussevskaia, W, FOWANC 2008
  - Hardness results for analog network coding
- Locher, von Rickenbach, W, ICDCN 2008
  - Still some major open problems

Main open question in this area

- Most papers so far deal with special cases, essentially scheduling a number of links with special properties. The general problem is still wide open:
- 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.
- E.g., for arbitrary power levels not even hardness is known…

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich