Wireless Algorithms

Roger Wattenhofer

… an oxymoron?

…as old as society

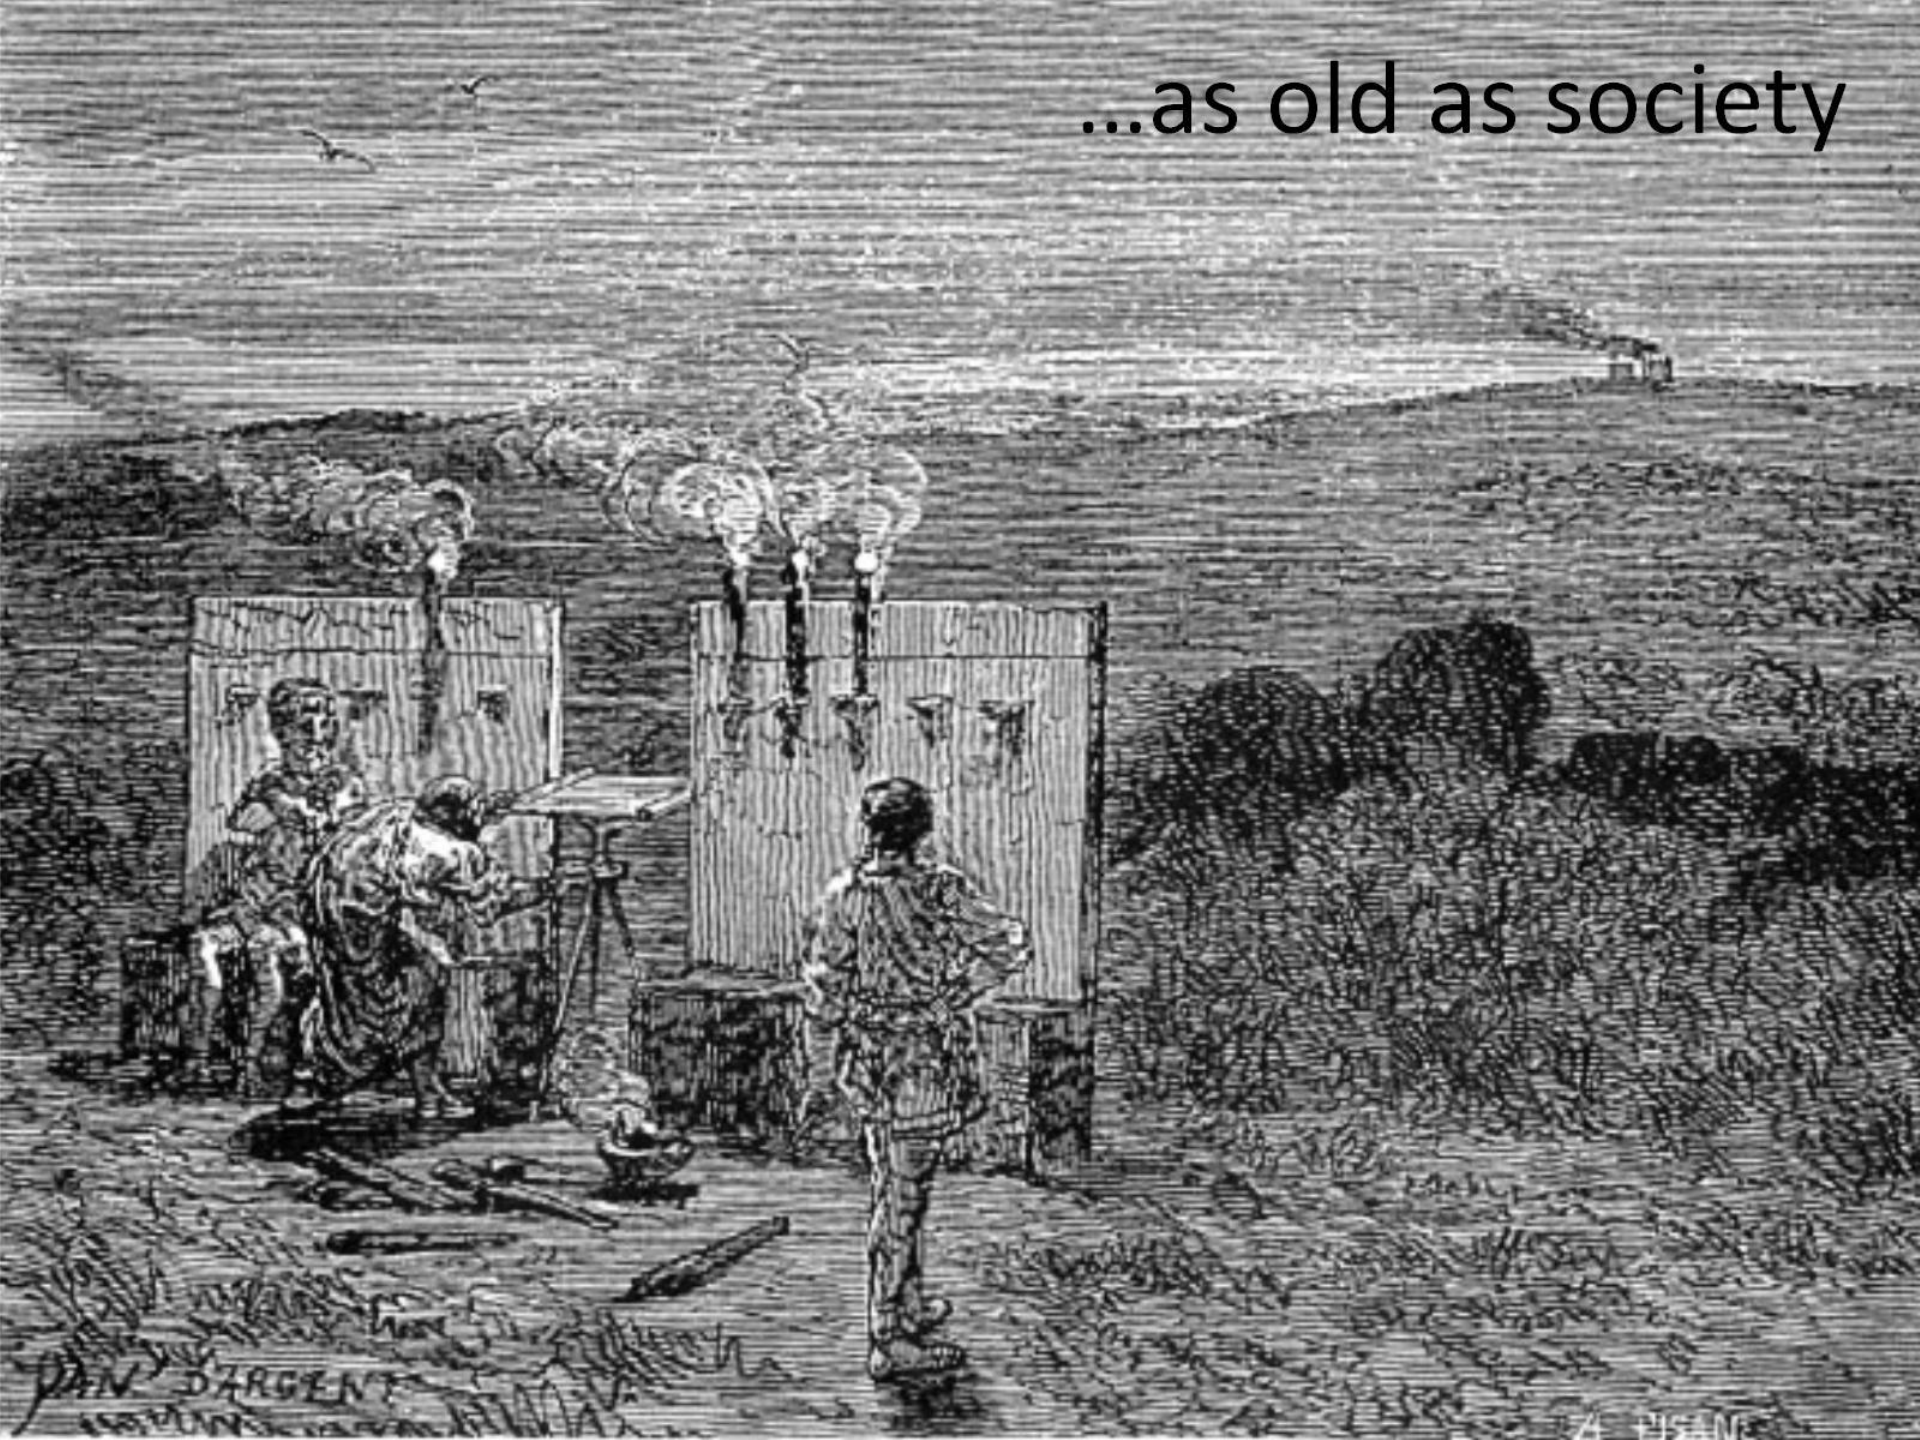

more recently?

Wireless Communication

EE, Physics

Maxwell Equations

Simulation, Testing

‘Scaling Laws’

Network Algorithms

CS, Applied Math

[Geometric] Graphs

Worst-Case Analysis

Any-Case Analysis

CS Models: e.g. Disk Model (Protocol Model)

CS Models: e.g. Disk Model (Protocol Model)

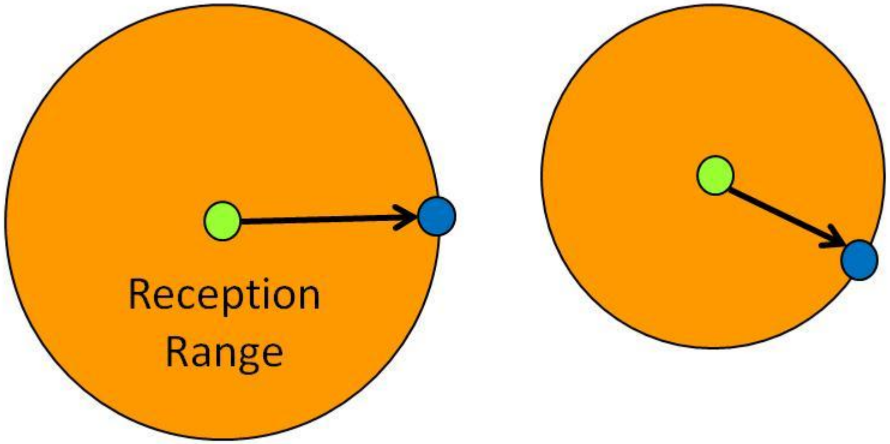

CS Models: e.g. Disk Model (Protocol Model)

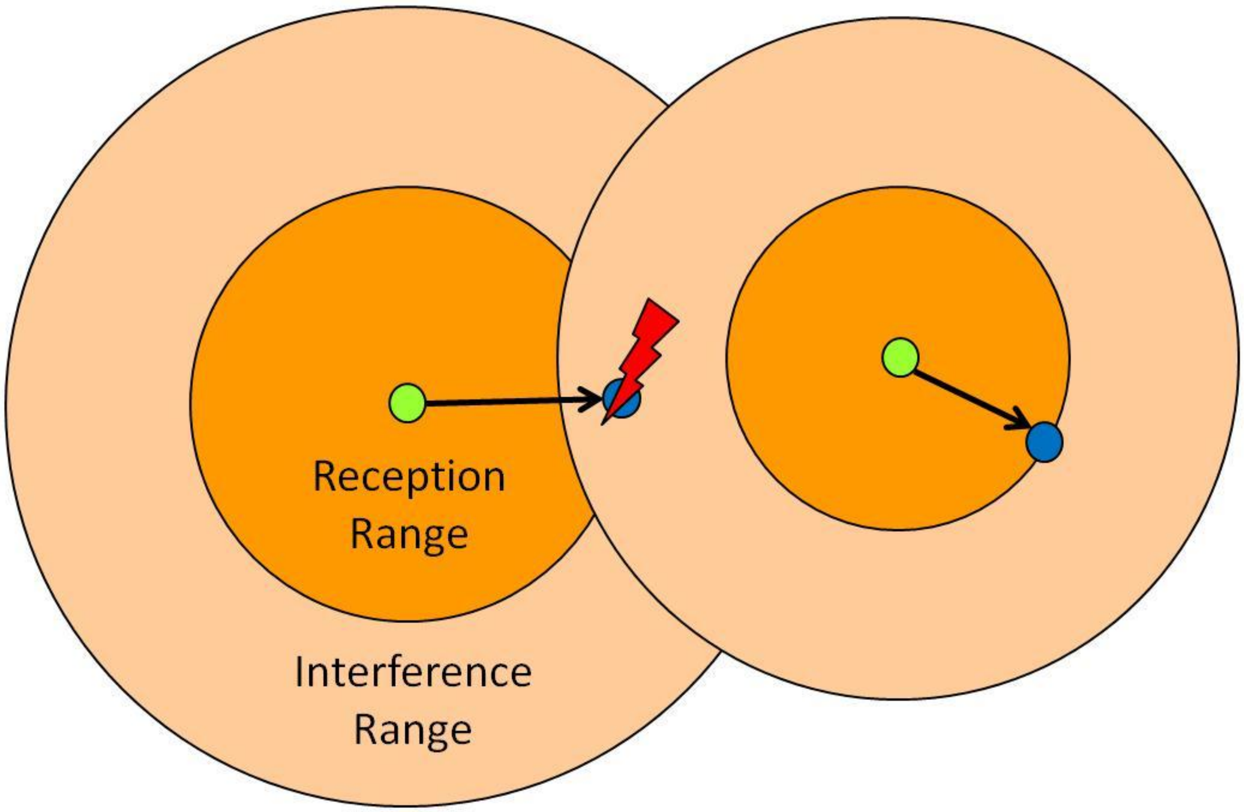

EE Models: e.g. SINR Model (Physical Model)

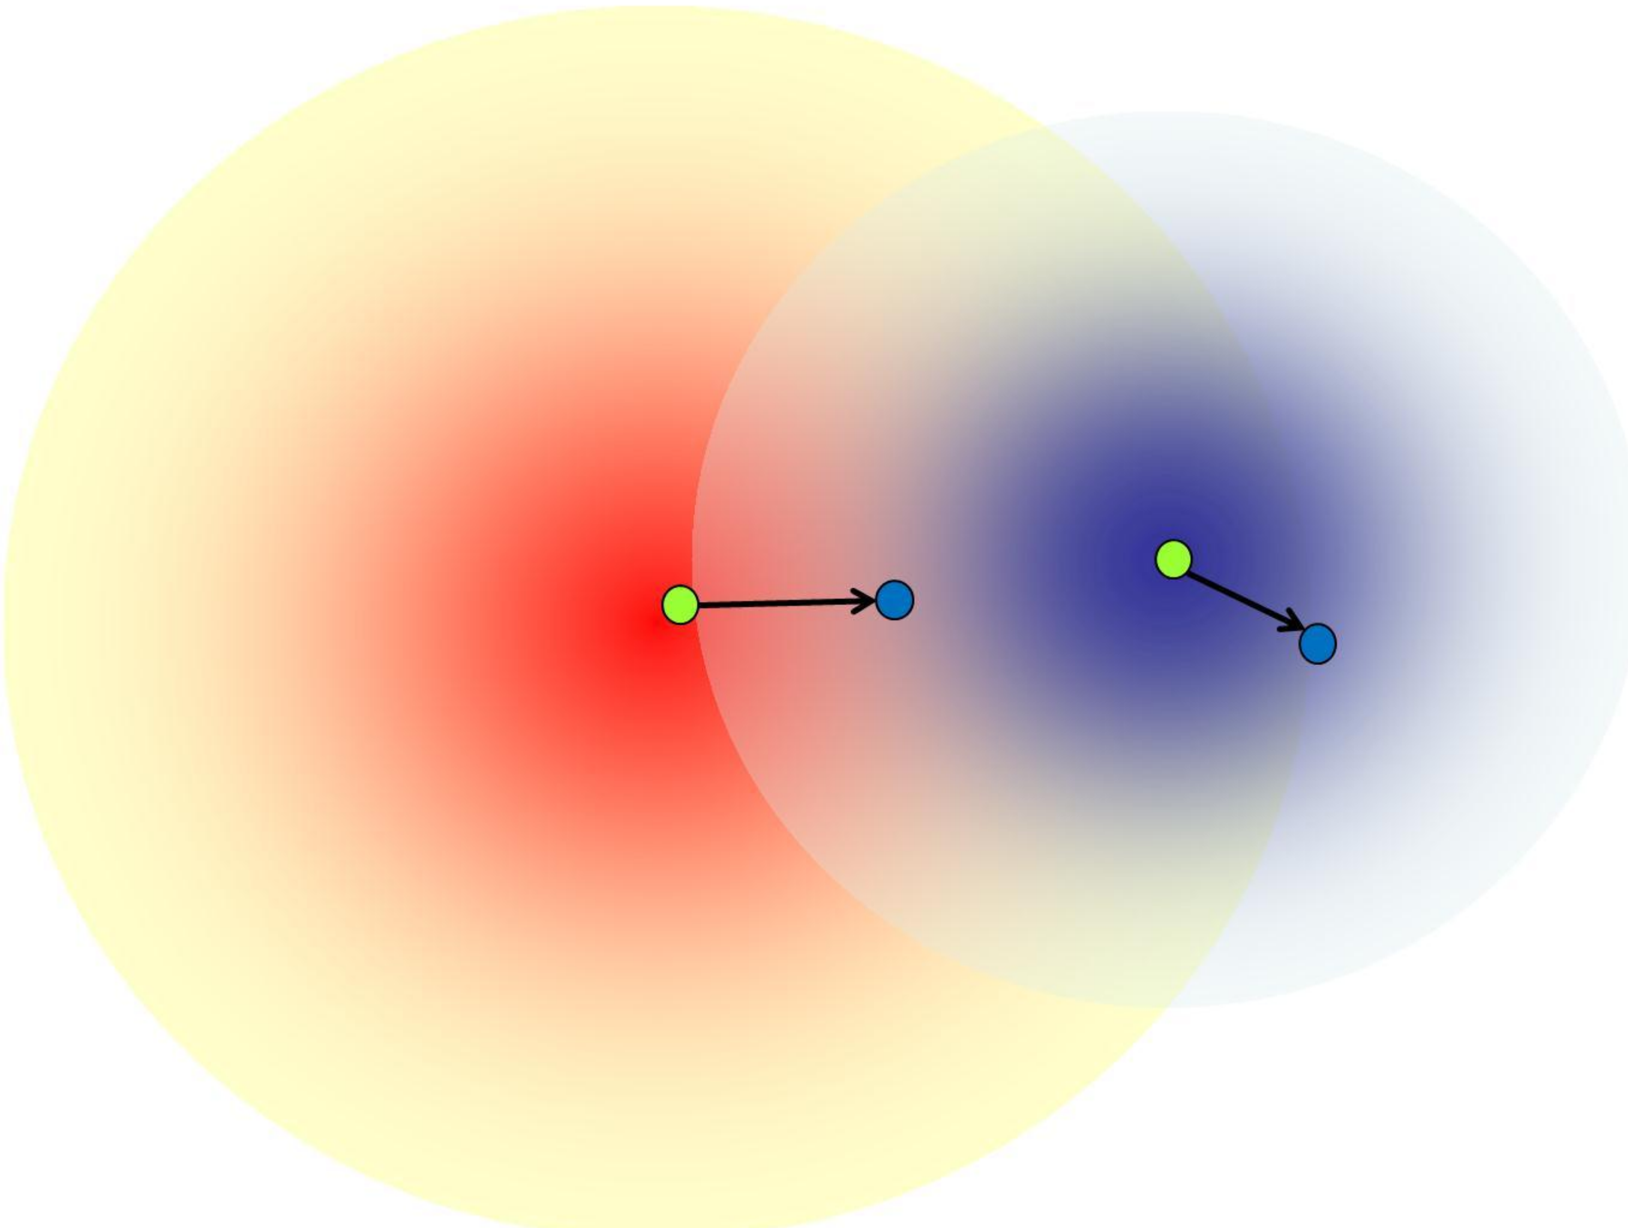

Signal-To-Interference-Plus-Noise Ratio (SINR) Formula

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

Signal-To-Interference-Plus-Noise Ratio (SINR) Formula

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

Received signal power from sender

Power level of sender *u*

Path-loss exponent

Minimum signal-to-interference ratio

Noise

Received signal power from all other nodes (=interference)

Distance between two nodes

Example: Protocol vs. Physical Model

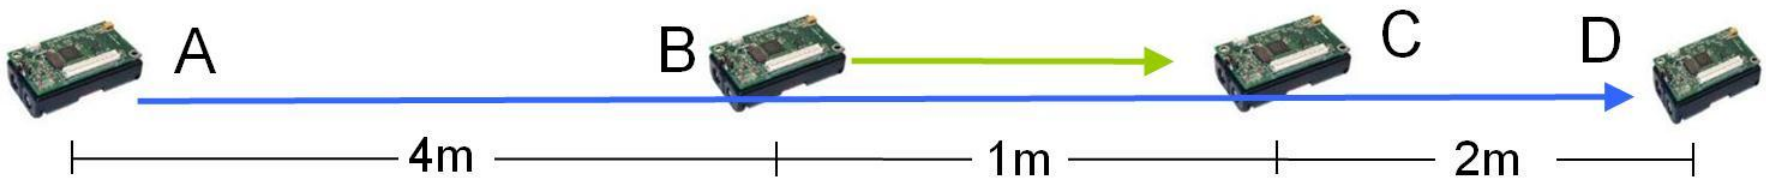

Assume a single frequency (and no fancy decoding techniques!)

Example: Protocol vs. Physical Model

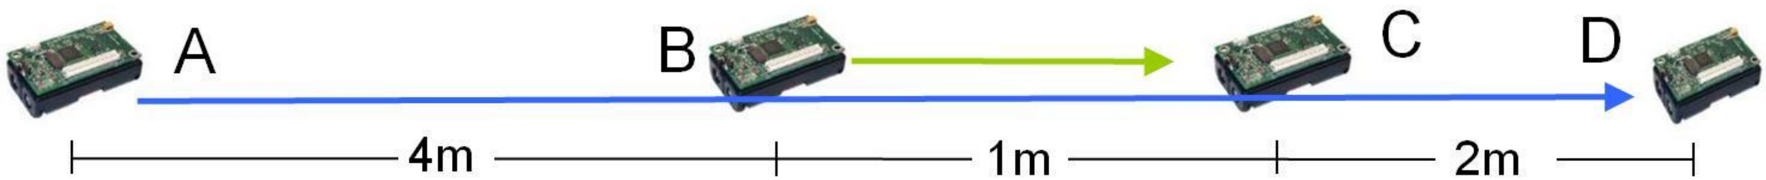

Assume a single frequency (and no fancy decoding techniques!)

Is spatial reuse possible? → NO — Protocol Model

Is spatial reuse possible? → YES — With power control

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!

… even with very simple hardware (sensor nodes)

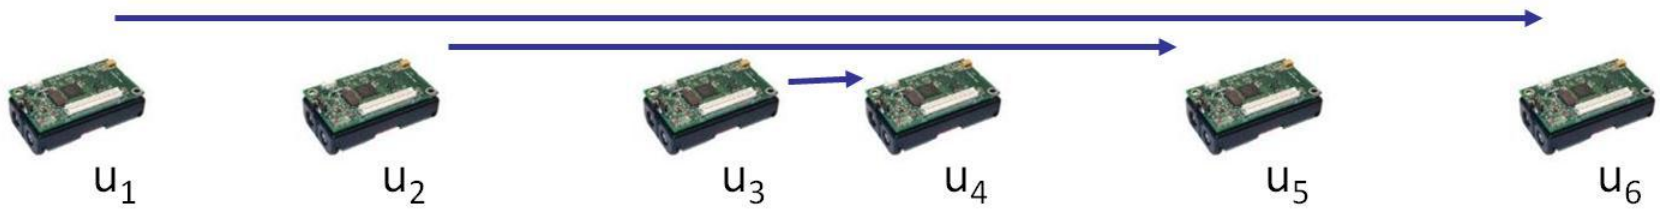

Time for transmitting 20‘000 packets:

	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, W, Weber, Hotnets 2006]

The Capacity of a Wireless Network

Measures for Capacity

Throughput capacity

- Number of packets successfully delivered per time
- Dependent on the traffic pattern
- *E.g.: What is the maximum achievable rate, over all protocols, for a random node distribution and a random destination for each source?*

Transport capacity

- A network transports one bit-meter when one bit has been transported a distance of one meter.
- *What is the maximum achievable rate, over all node locations, and all traffic patterns, and all protocols?*

Convergecast capacity

- *How long does it take to get information from all nodes to a sink*

Many more…

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.
- Quiz: 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

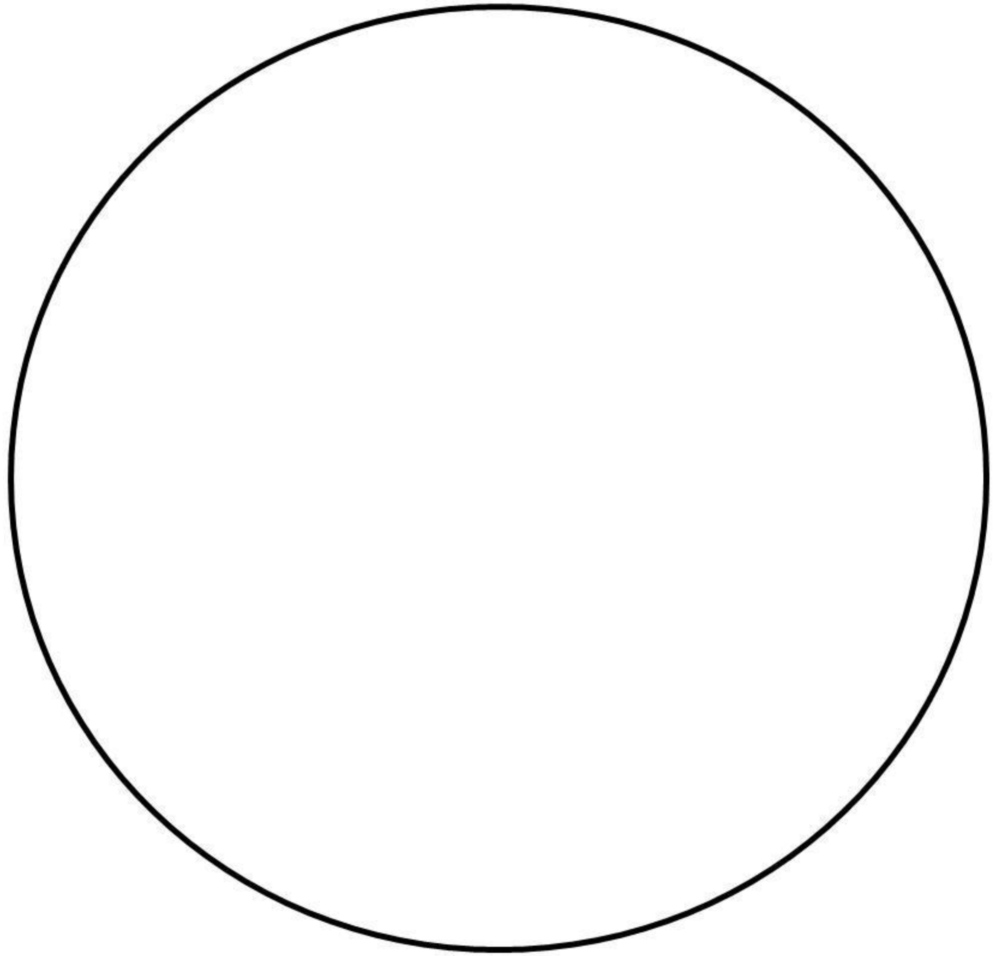

Transport Capacity: Example

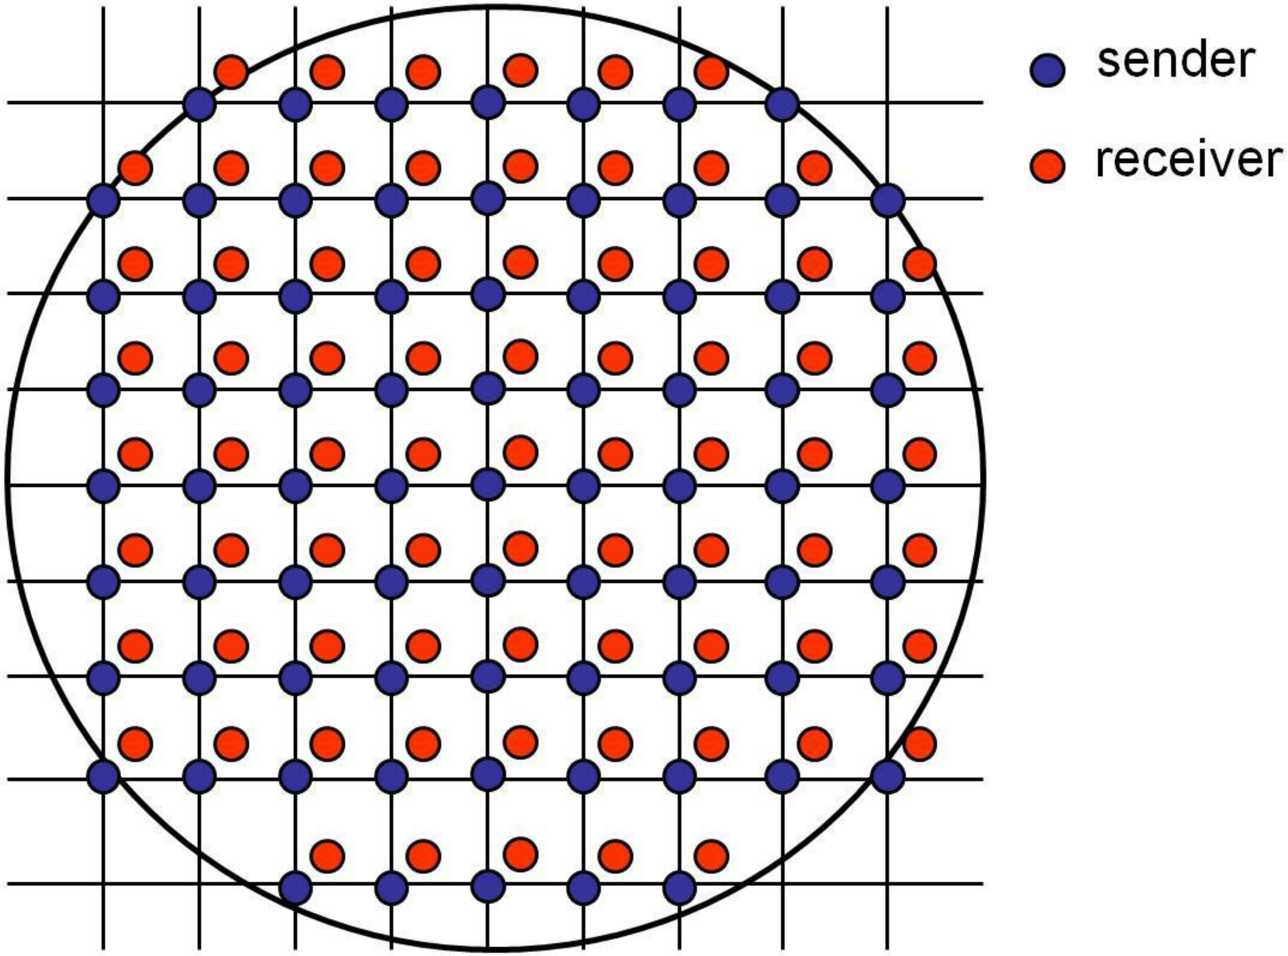

Transport Capacity: Understanding the example

- 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? No, 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 $\Theta(d^2)$.

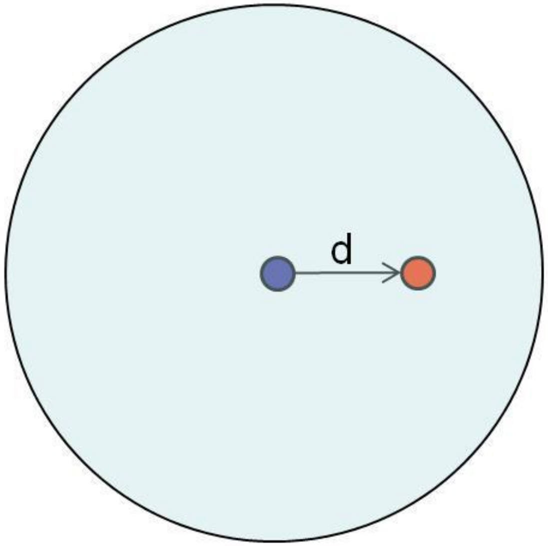

- We want to maximize n transmissions with distances $d_1, d_2, \ldots, d_n$ given that the total area is less than a unit disk. This is maximized if all $d_i = \Theta(1/\sqrt{n})$. So the example was asymptotically optimal.
  - BTW, a fun open geometry problem: Given k circles with total area 1, can you always fit them in a circle with total area 2?

More capacities…

- The throughput capacity of an *n* node random network is $\Theta(\frac{W}{\sqrt{n \log n}})$
- There exist constants c and c' such that $\lim_{n\to\infty} \Pr[c\frac{W}{\sqrt{n \log n}} \text{ is feasible}] = 1$

$$\lim_{n\to\infty} \Pr[c'\frac{W}{\sqrt{n \log n}} \text{ is feasible}] = 0$$

- 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…

- There are literally thousands of results on capacity, e.g.
  - on random destinations
  - on power-law traffic patterns
  - communication through relays
  - communication in mobile networks
  - channel broken into subchannels
  - etc.

Practical relevance?

- Efficient access to media, i.e. MAC layer!
- This (and related) problem is studied theoretically:

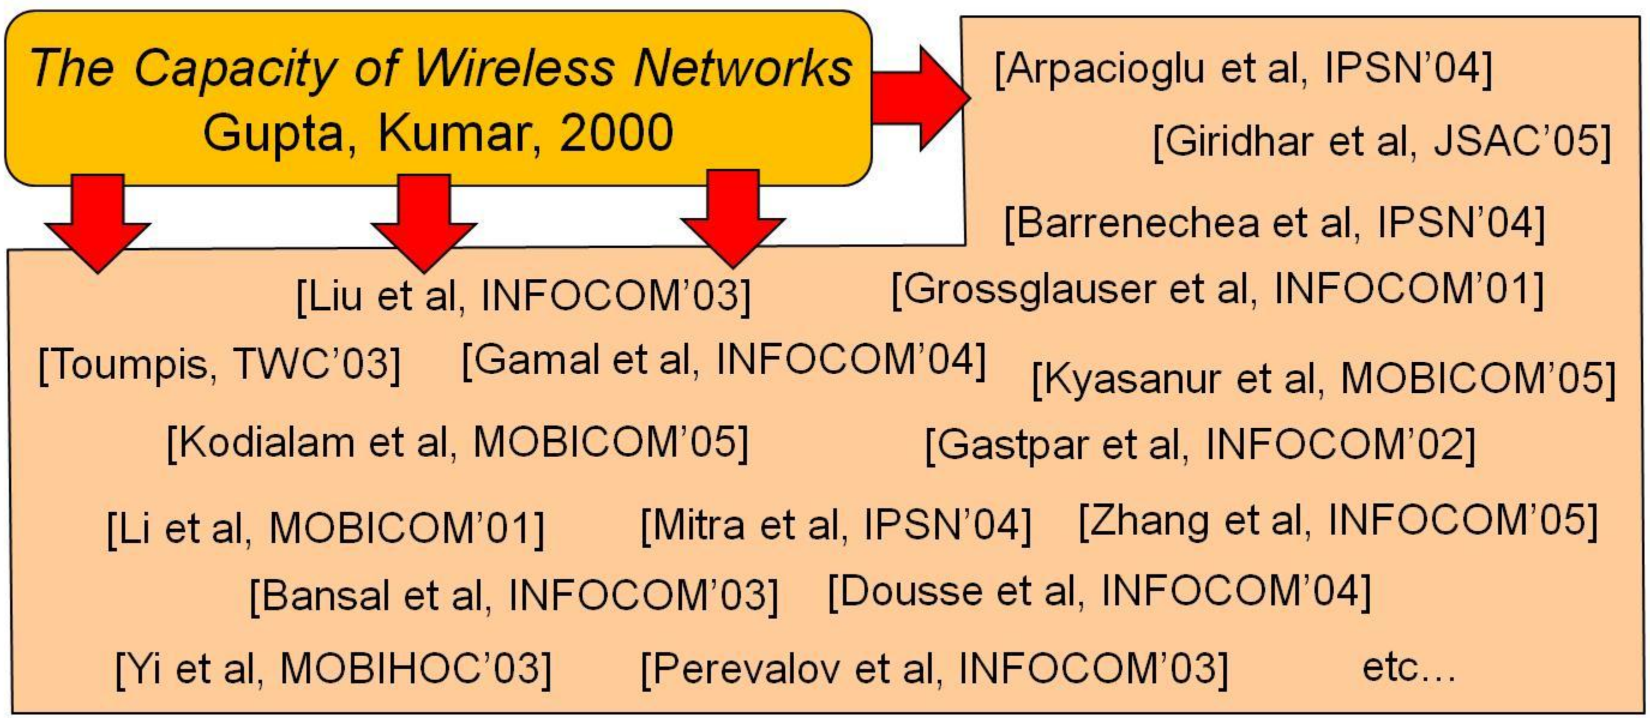

Network Topology?

- All these capacity studies make very strong assumptions on node deployment, topologies
  - randomly, uniformly distributed nodes
  - nodes placed on a grid
  - etc.

What if a network looks differently…?

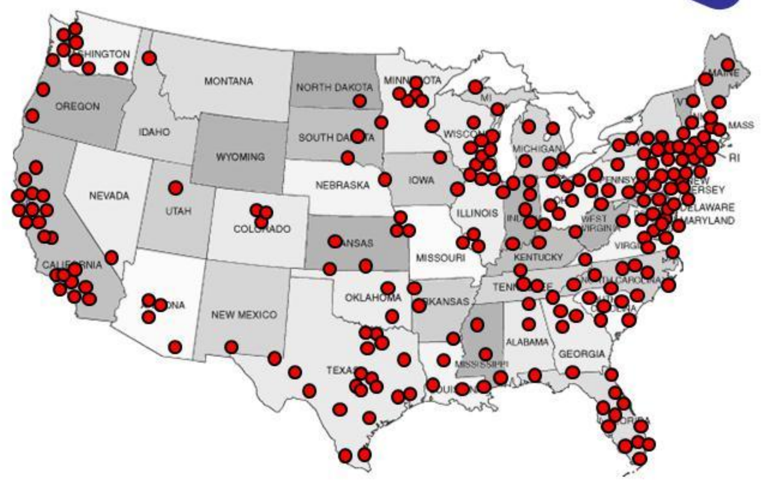

‘Scaling Laws’

“Classic” Capacity

How much information can be transmitted in nice networks?

'Scaling Laws'

"Classic" Capacity

How much information can be transmitted in nice networks?

Worst-Case Capacity

How much information can be transmitted in nasty networks?

‘Scaling Laws’

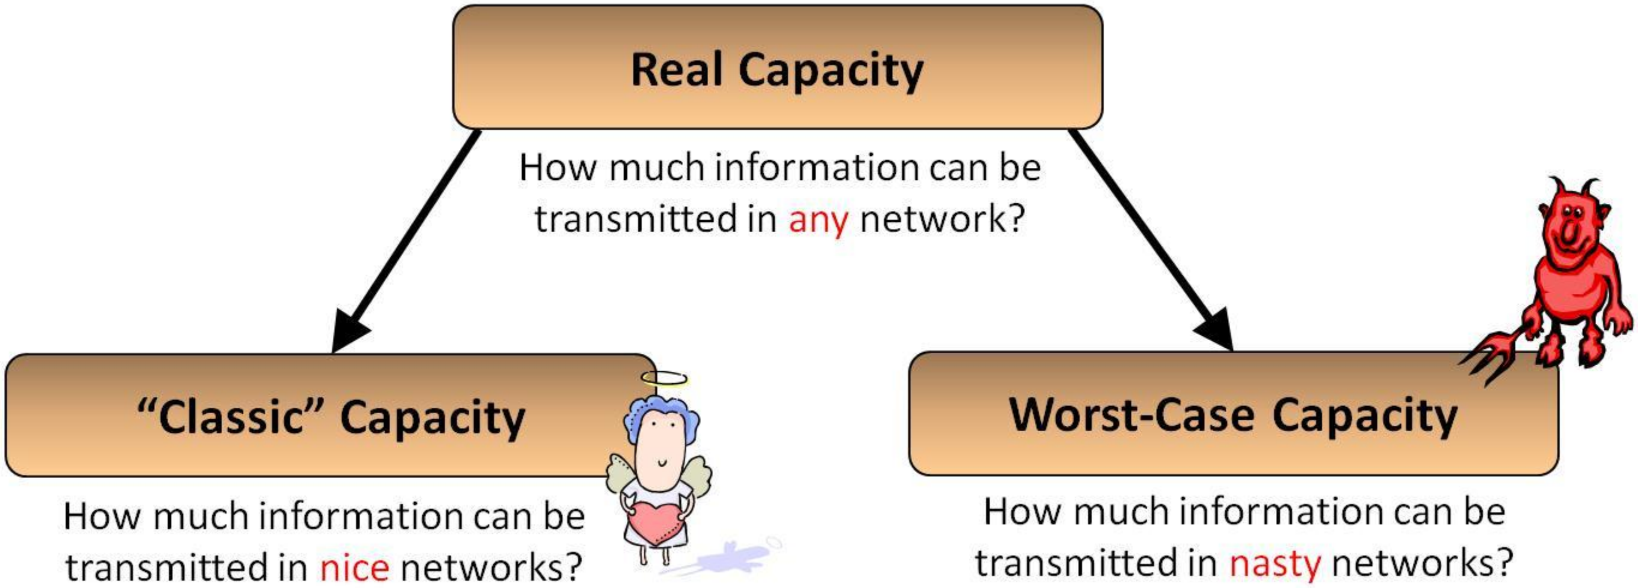

Convergecast Capacity in Wireless Networks

- 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?

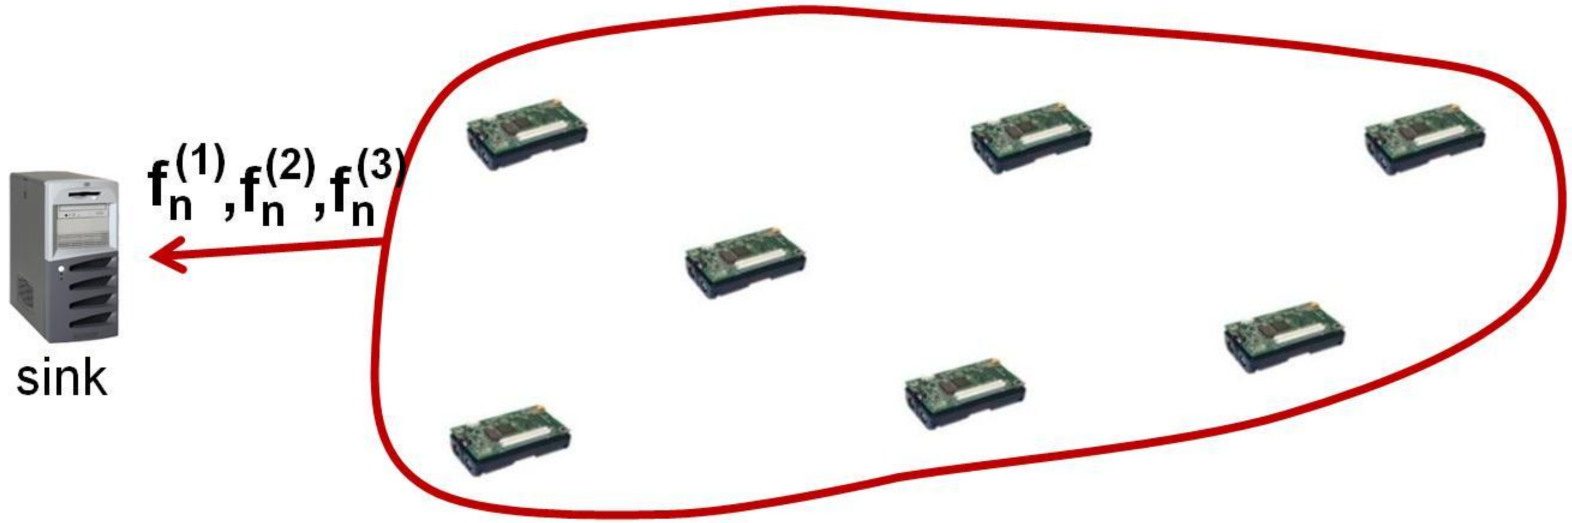

Convergecast Capacity in Wireless Networks

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

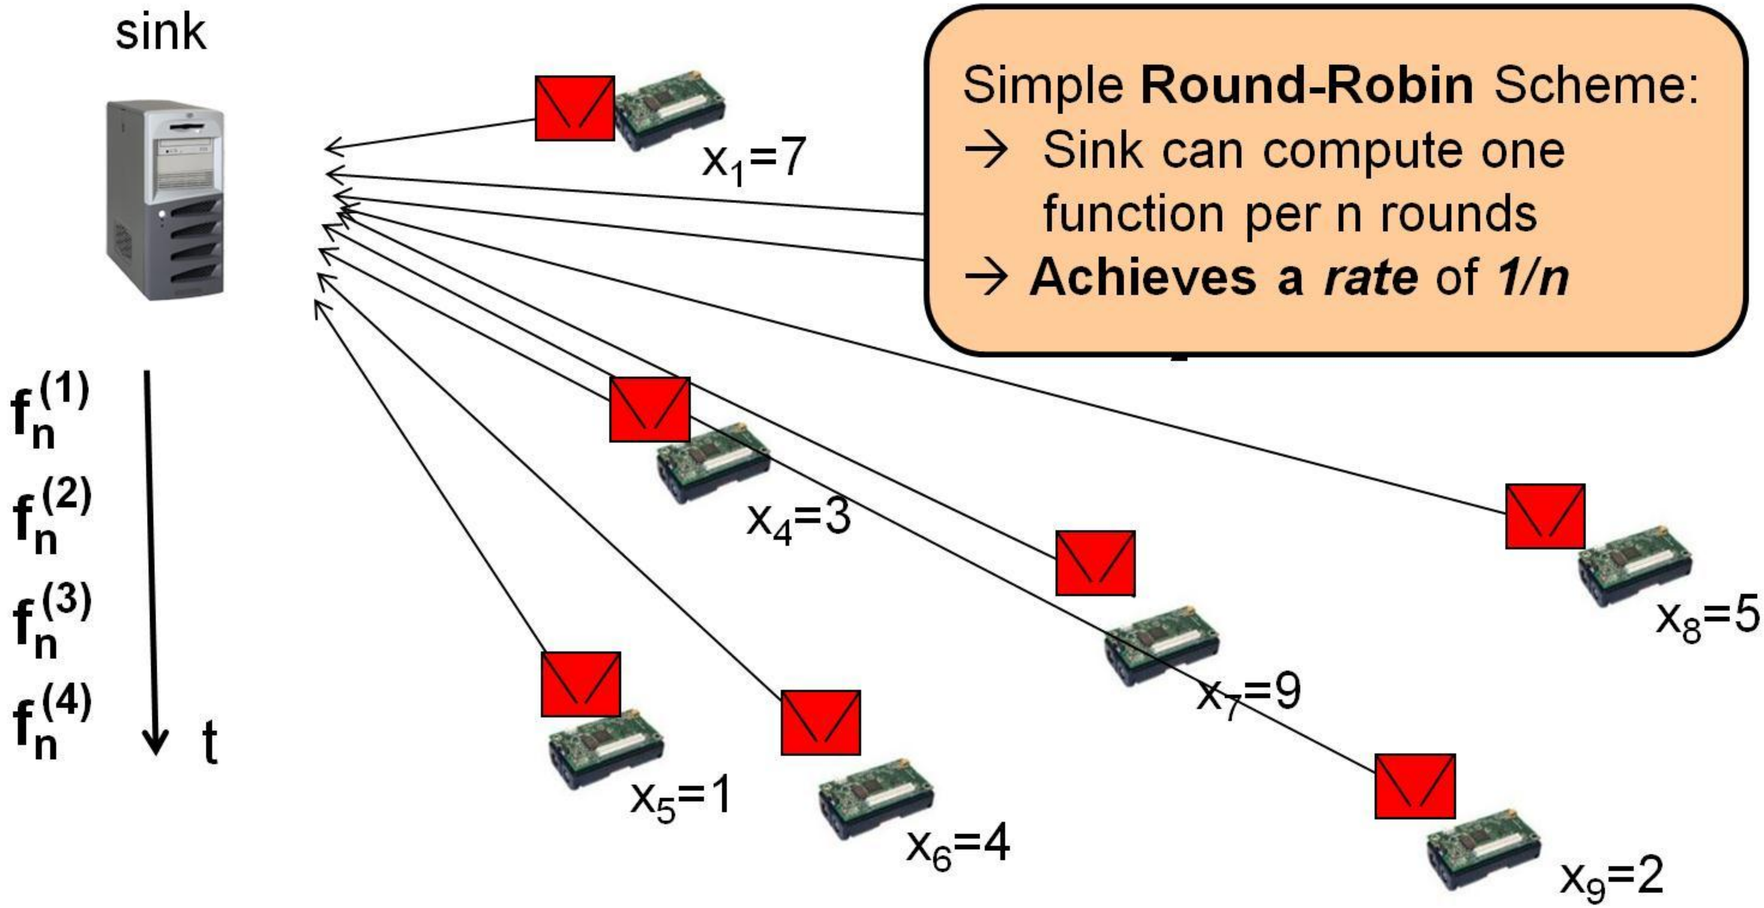

Simple **Round-Robin** Scheme:
→ Sink can compute one function per n rounds
→ **Achieves a *rate* of *1/n***

Convergecast Capacity in Wireless Networks

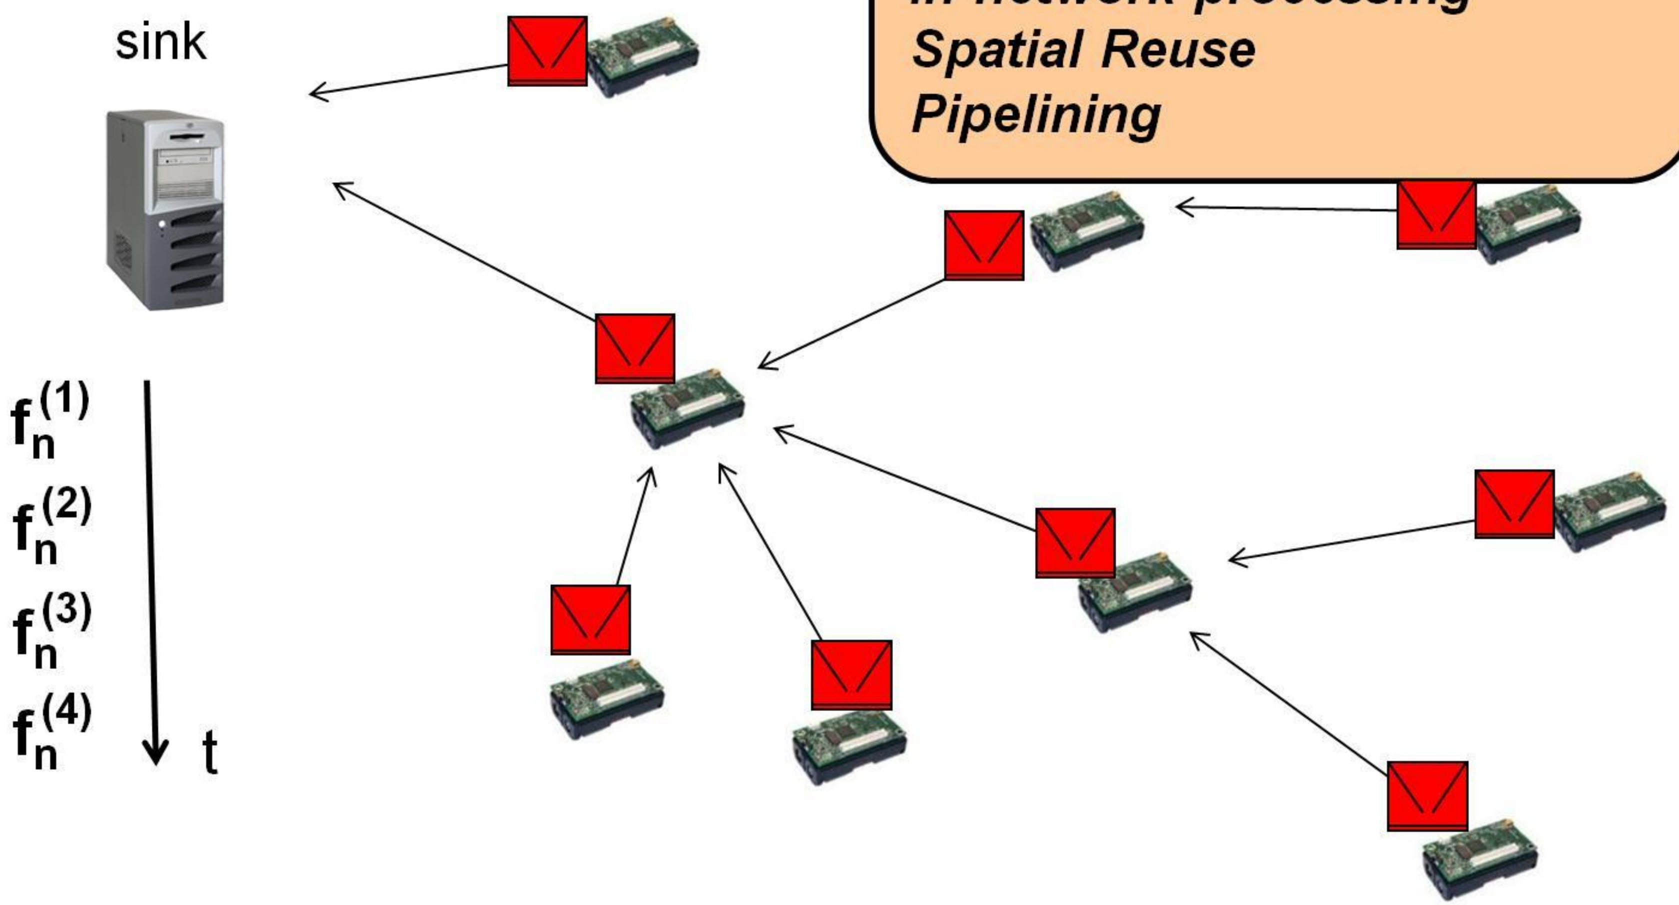

There are better schemes using
Multi-hop relaying
In-network processing
Spatial Reuse
Pipelining

Convergecast Capacity in Wireless Networks

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

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

Convergecast Capacity in Wireless 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:

- Function to be computed → Only perfectly compressible functions (max, min, avg,…)
- Coding techniques → No fancy coding techniques
- Network topology
- …

Convergecast Capacity in Wireless Networks

Networks / Model/Power	Max. rate in arbitrary, worst-case deployment	Max. rate in random, uniform deployment
Protocol Model		
Physical Model (w/ power control)		

Convergecast Capacity in Wireless Networks

[Moscibroda, W, 2006]

Worst-Case Capacity

[Giridhar, Kumar, 2005]

Best-Case Capacity

Networks / Model/Power	Max. rate in arbitrary, worst-case deployment	Max. rate in random, uniform deployment
Protocol Model		
Physical Model (w/ power control)		

Convergecast Capacity in Wireless Networks

[Moscibroda, W, 2006]

Worst-Case Capacity

[Giridhar, Kumar, 2005]

Best-Case Capacity

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

Convergecast Capacity in Wireless Networks

[Moscibroda, W, 2006]

Worst-Case Capacity

[Giridhar, Kumar, 2005]

Best-Case Capacity

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

Convergecast Capacity in Wireless Networks

[Moscibroda, W, 2006] — Worst-Case Capacity

[Giridhar, Kumar, 2005] — Best-Case Capacity

Model/Power \ 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 (w/ power control)	$\Omega(1/\log^3 n)$	$\Omega(1/\log n)$

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!)

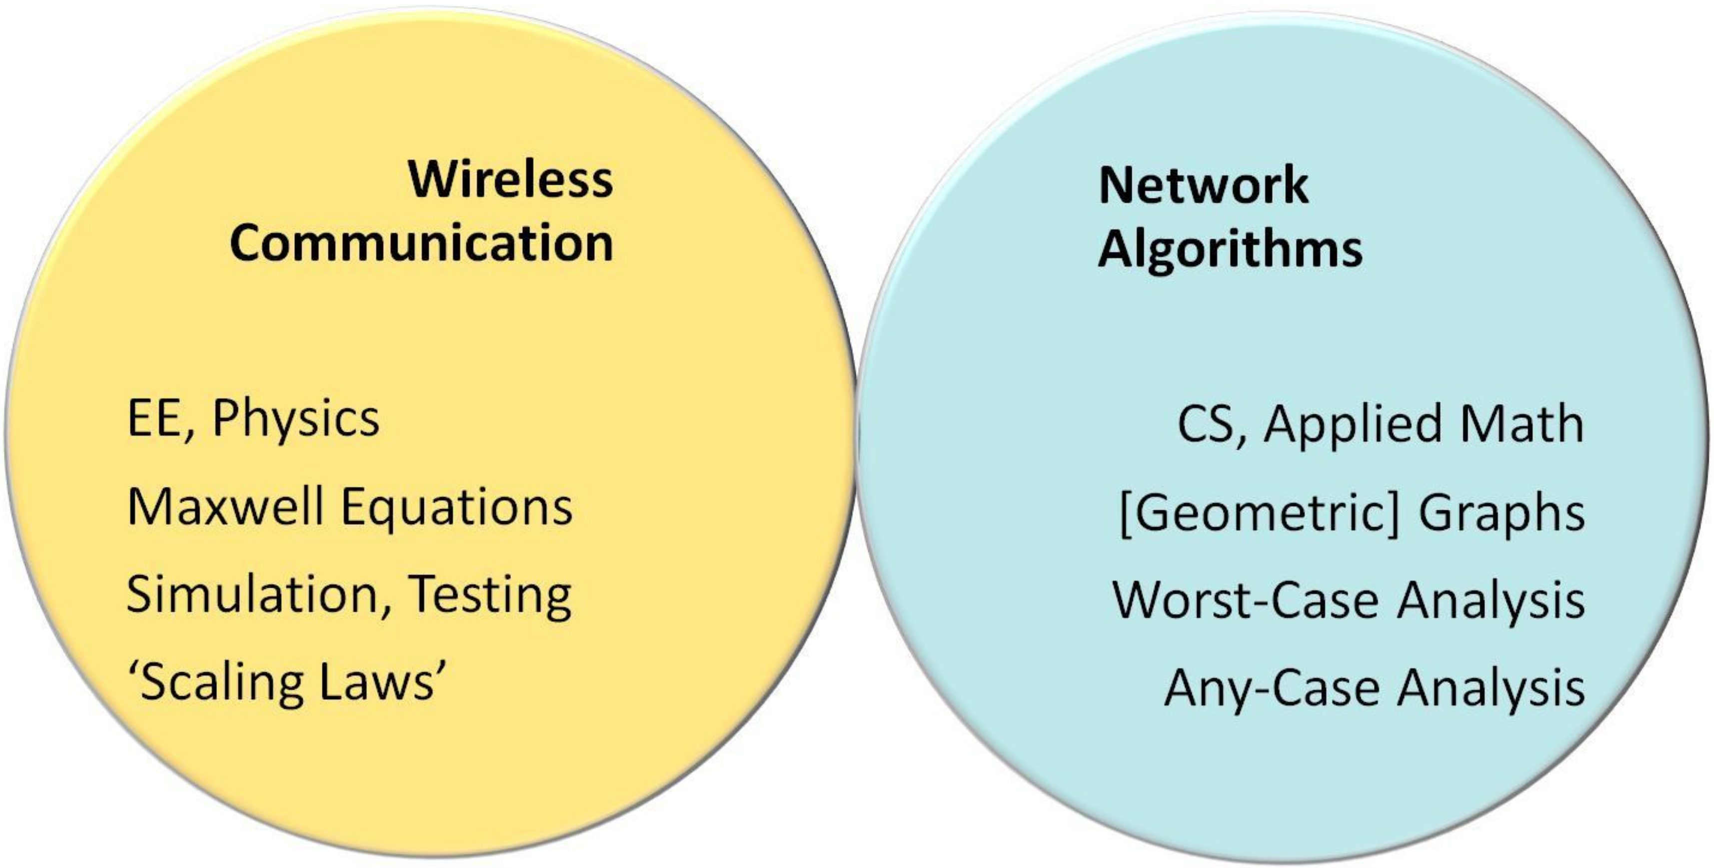
Wireless Communication
EE, Physics
Maxwell Equations
Simulation, Testing
'Scaling Laws'
Network Algorithms
CS, Applied Math
[Geometric] Graphs
Worst-Case Analysis
Any-Case Analysis

Possible Application – Hotspots in WLAN

Traditionally: clients assigned to (more or less) closest access point

→ far-terminal problem → hotspots have less throughput

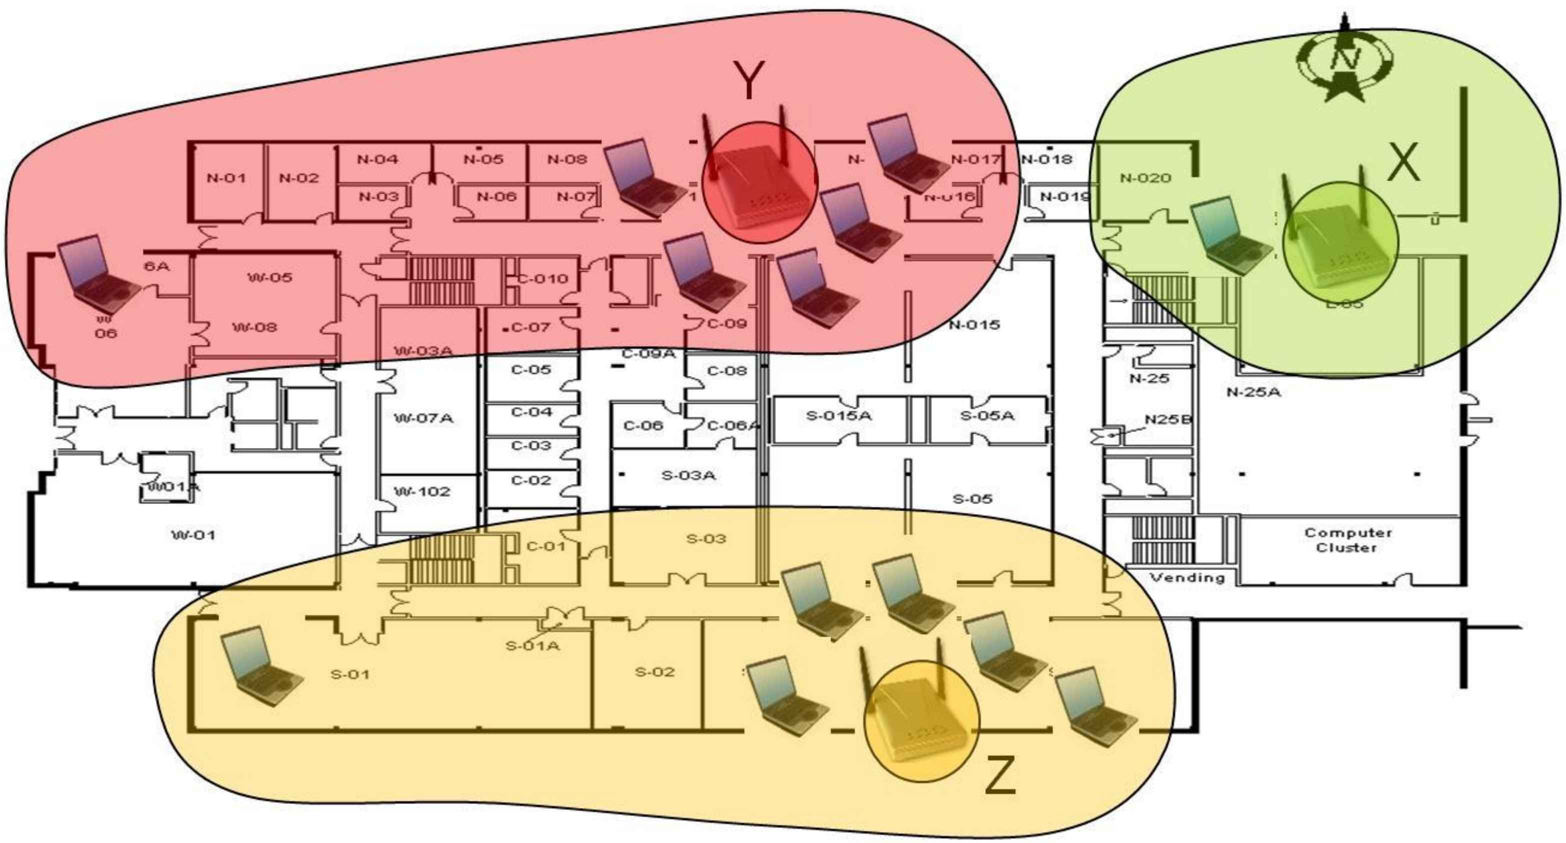

Possible Application – Hotspots in WLAN

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!

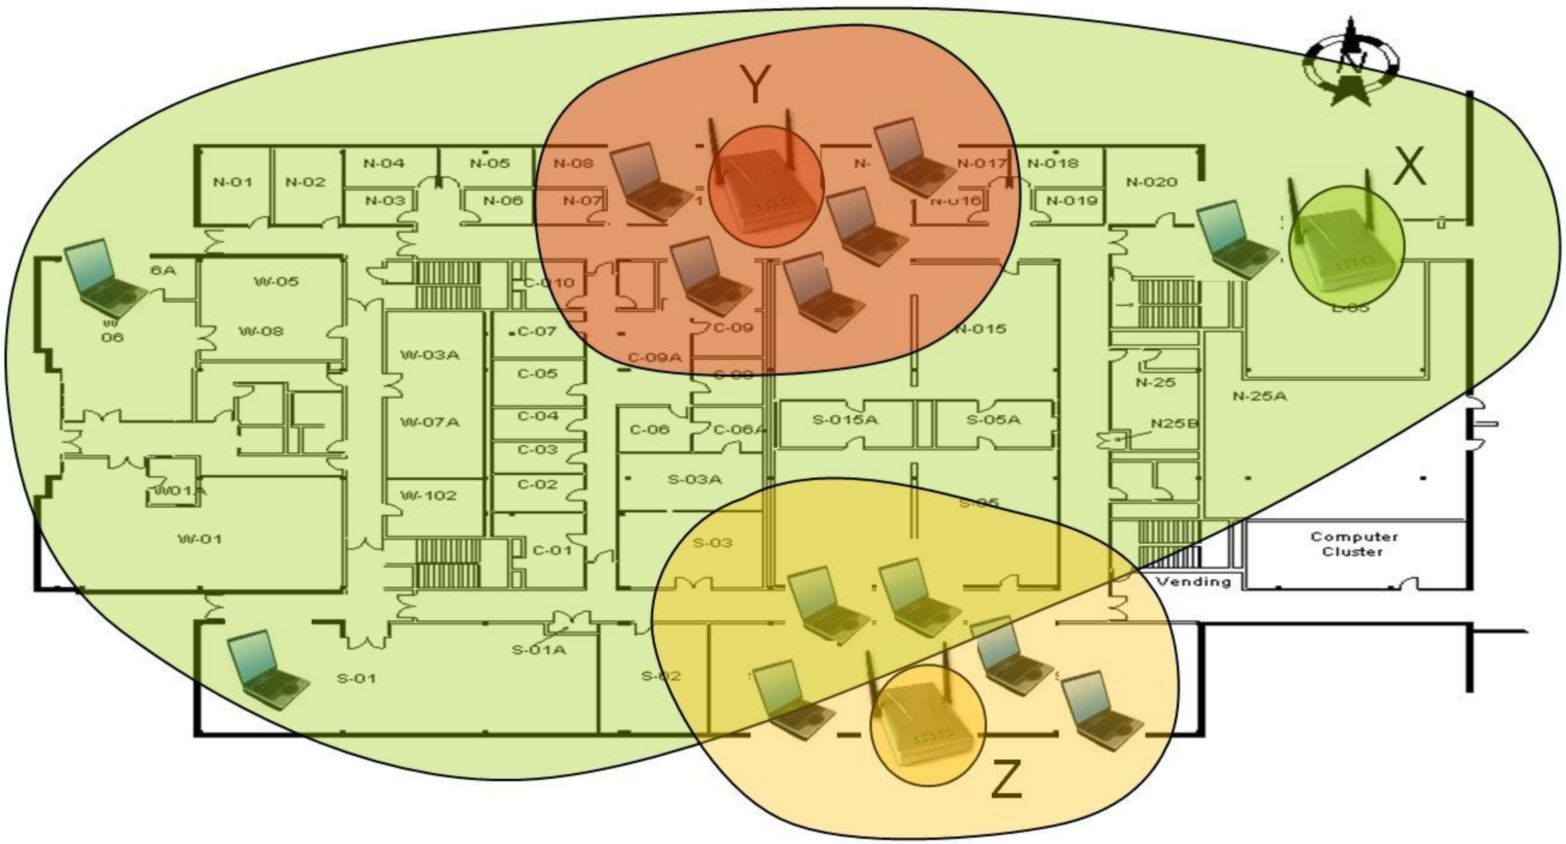

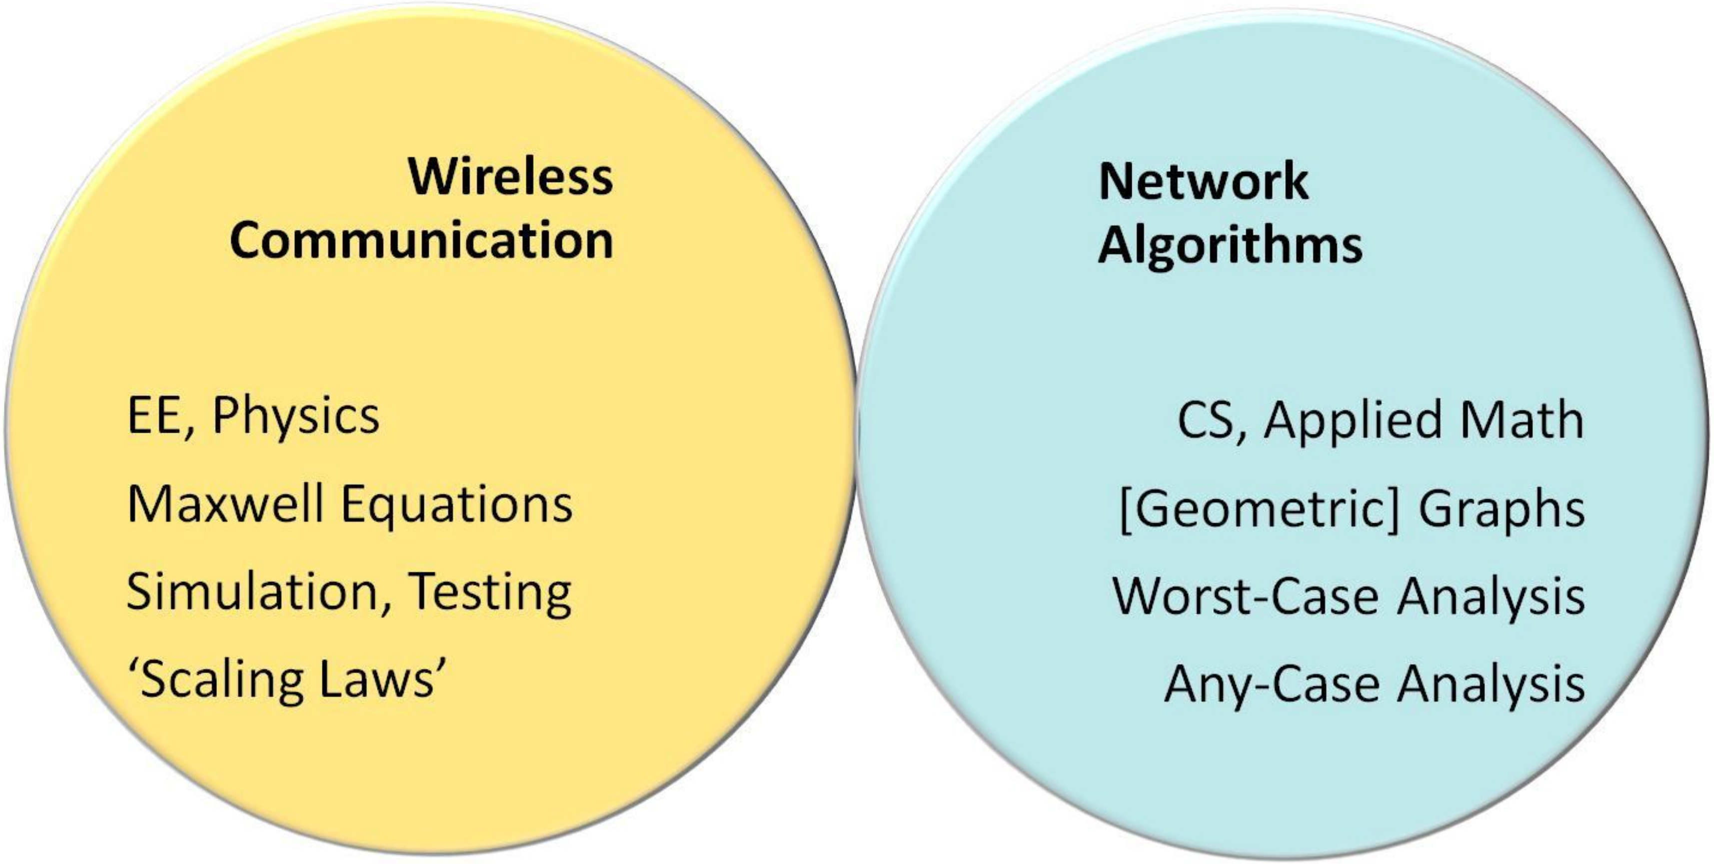
Wireless Communication
EE, Physics
Maxwell Equations
Simulation, Testing
'Scaling Laws'
Network Algorithms
CS, Applied Math
[Geometric] Graphs
Worst-Case Analysis
Any-Case Analysis

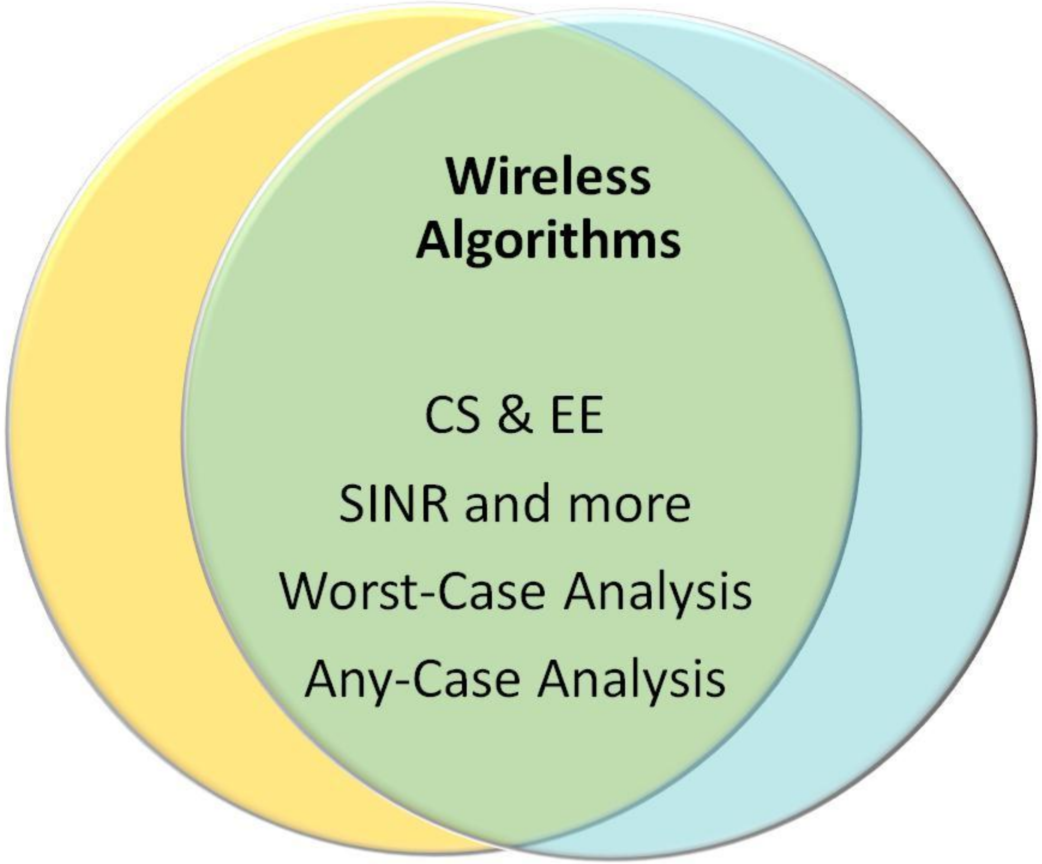

Wireless Algorithms
CS & EE
SINR and more
Worst-Case Analysis
Any-Case Analysis

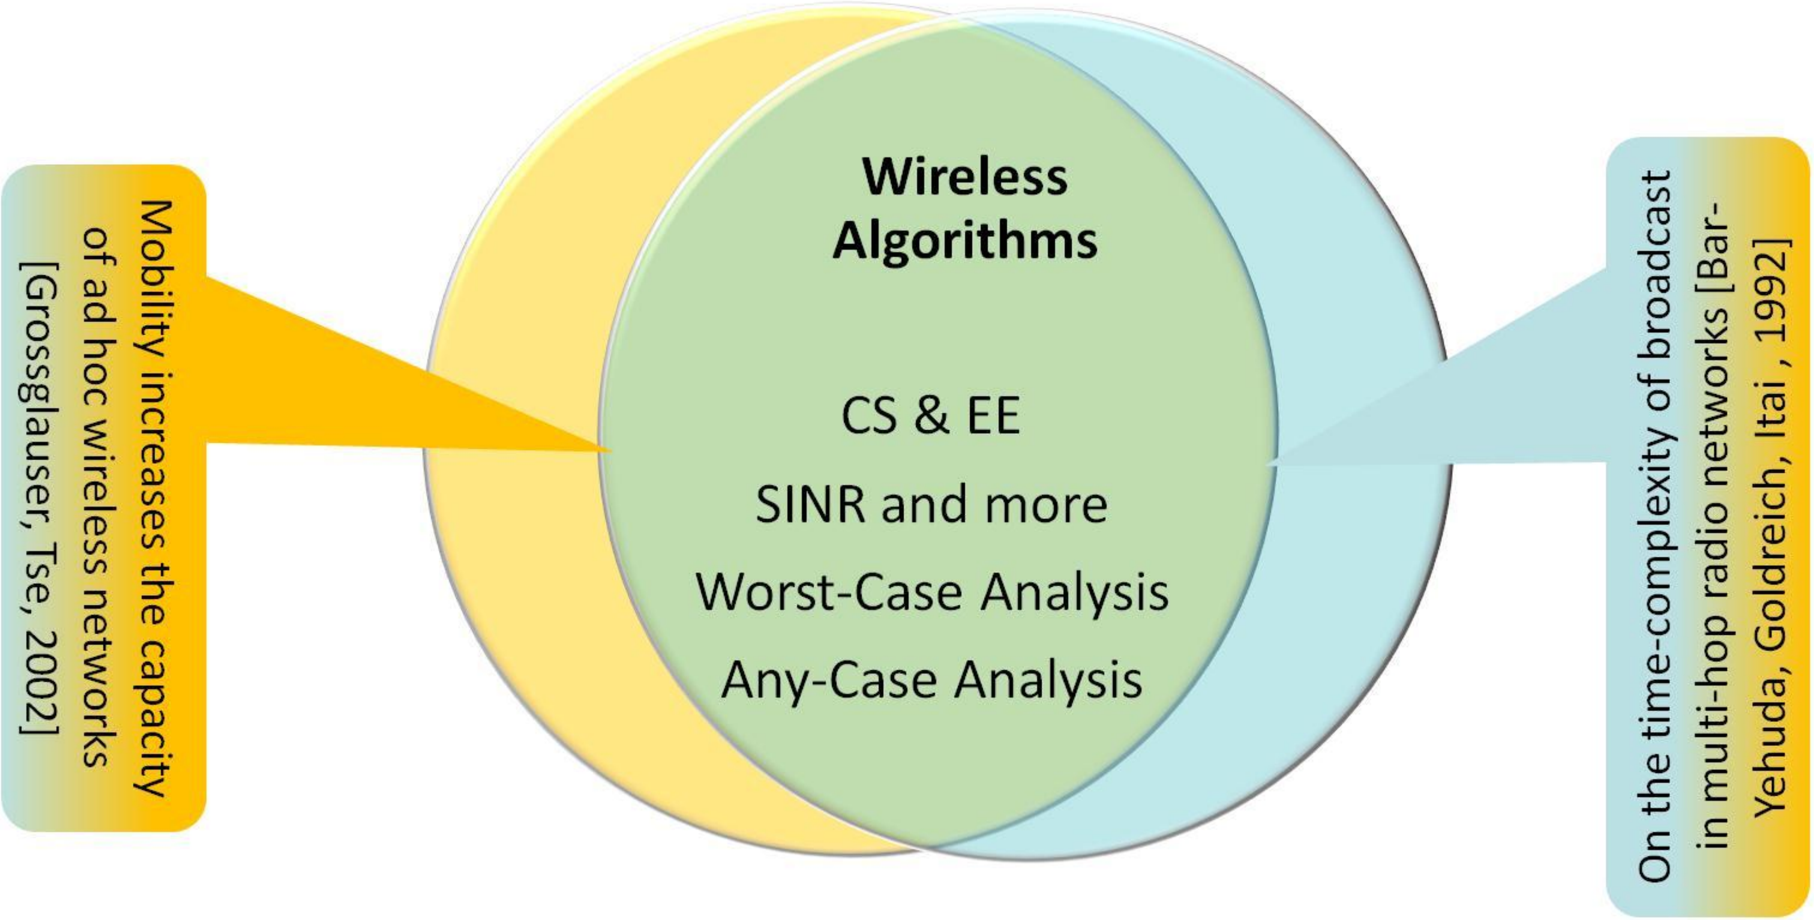

Mobility increases the capacity
of ad hoc wireless networks
[Grossglauser, Tse, 2002]
Wireless
Algorithms
CS & EE
SINR and more
Worst-Case Analysis
Any-Case Analysis
On the time-complexity of broadcast
in multi-hop radio networks [Bar-
Yehuda, Goldreich, Itai , 1992]

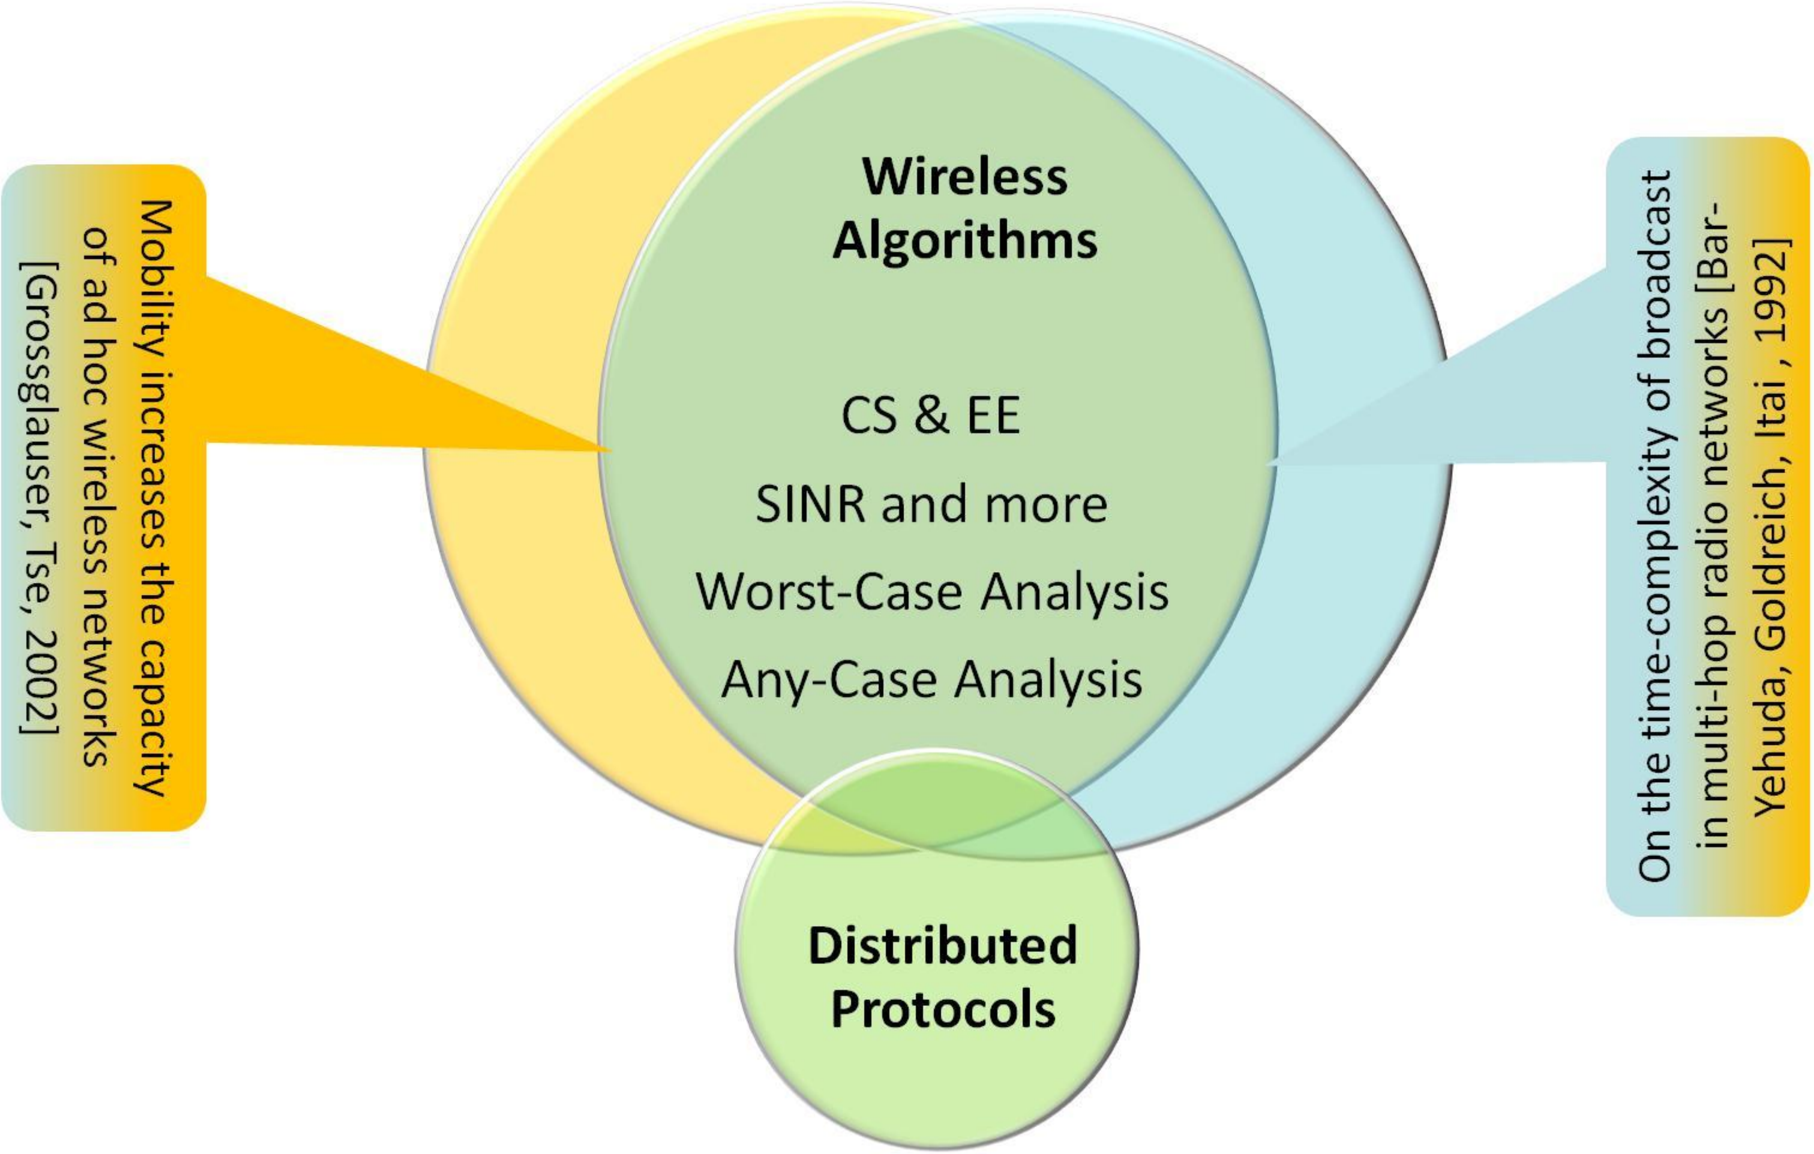
Mobility increases the capacity
of ad hoc wireless networks
[Grossglauser, Tse, 2002]
Wireless
Algorithms
CS & EE
SINR and more
Worst-Case Analysis
Any-Case Analysis
Distributed
Protocols
On the time-complexity of broadcast
in multi-hop radio networks [Bar-
Yehuda, Goldreich, Itai , 1992]

Wireless Communication 101

Signal-To-Interference-Plus-Noise Ratio (SINR) Formula

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

Received signal power from sender

Power level of sender *u*

Path-loss exponent

Noise

Minimum signal-to-interference ratio

Received signal power from all other nodes (=interference)

Dist etween

des

Ratio β depends on receiver (hardware, software, parameters)

- Simple solutions have $\beta > 10$
  - But $\beta < 1$ is possible (thanks to forward error correction)

Ratio β depends on receiver (hardware, software, parameters)

- Simple solutions have $\beta > 10$
  - But $\beta < 1$ is possible (thanks to forward error correction)

- Algorithmically speaking, the exact value of β does not really matter, thanks to SINR robustness
  - [Halldorsson, W, 2009] and [Fanghänel, Kesselheim, Räcke, Vöcking, 2009]
  - Model not only robust with regard to β, but also with regard to other constant factor disturbances, for instance, wind, constant antenna gain, etc.
  - Concretely: If we adapt model by factor ϕ, results will change at most by factor ϕ^2.

Ratio β depends on receiver (hardware, software, parameters)

- Simple solutions have $\beta > 10$
  - But $\beta < 1$ is possible (thanks to forward error correction)

- Algorithmically speaking, the exact value of β does not really matter, thanks to SINR robustness
  - [Halldorsson, W, 2009] and [Fanghänel, Kesselheim, Räcke, Vöcking, 2009]
  - Model not only robust with regard to β, but also with regard to other constant factor disturbances, for instance, wind, constant antenna gain, etc.
  - Concretely: If we adapt model by factor ϕ, results will change at most by factor ϕ^2.

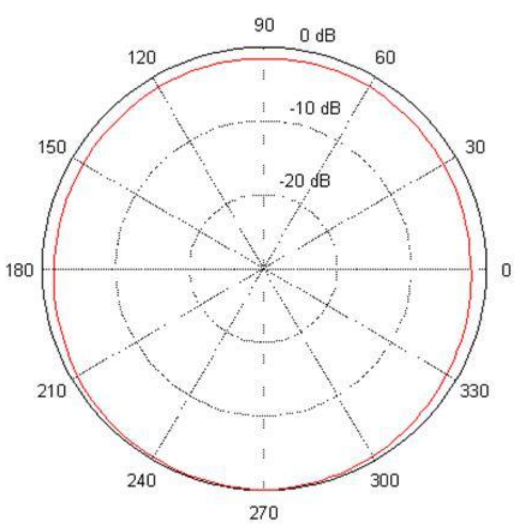

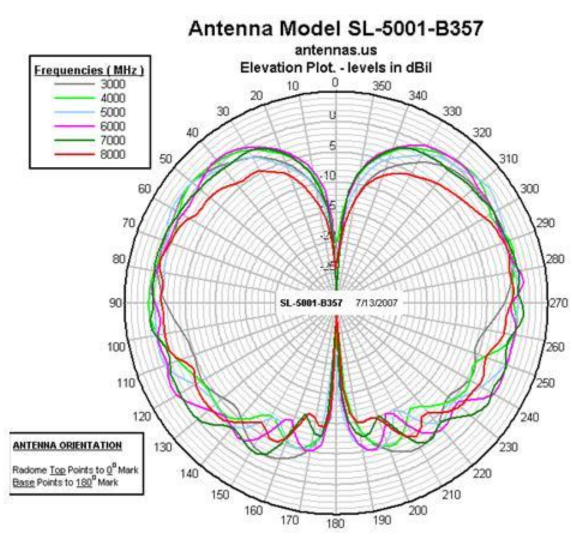

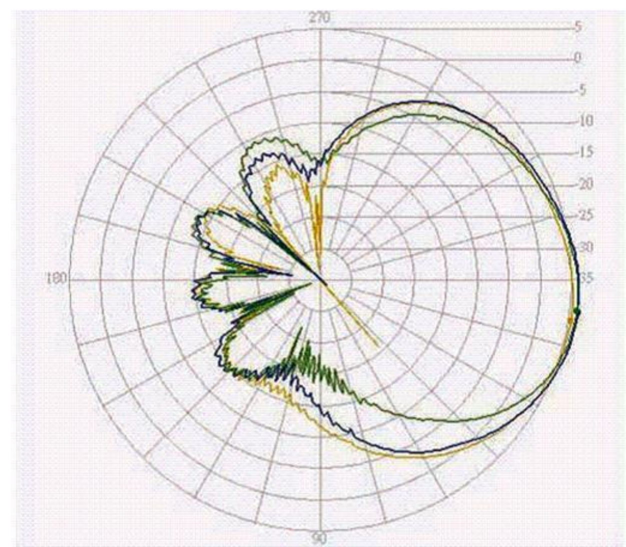

Modulation and demodulation

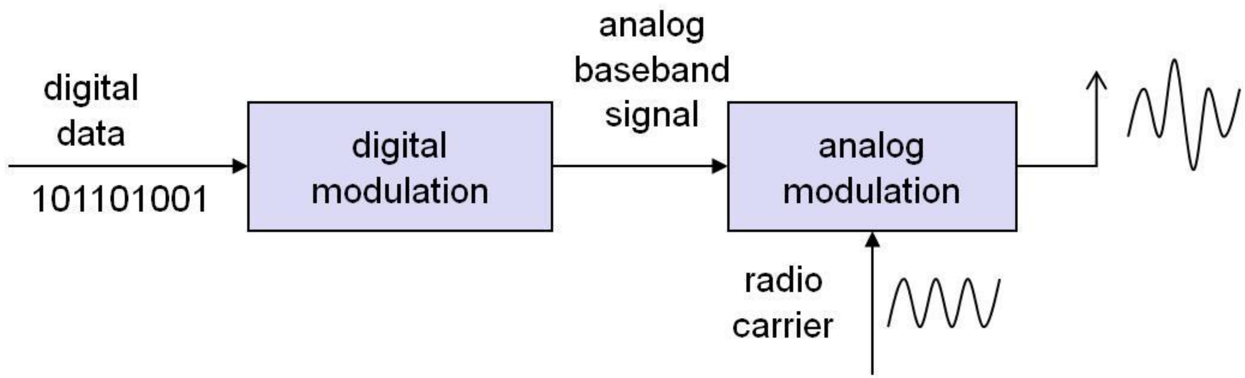

radio transmitter

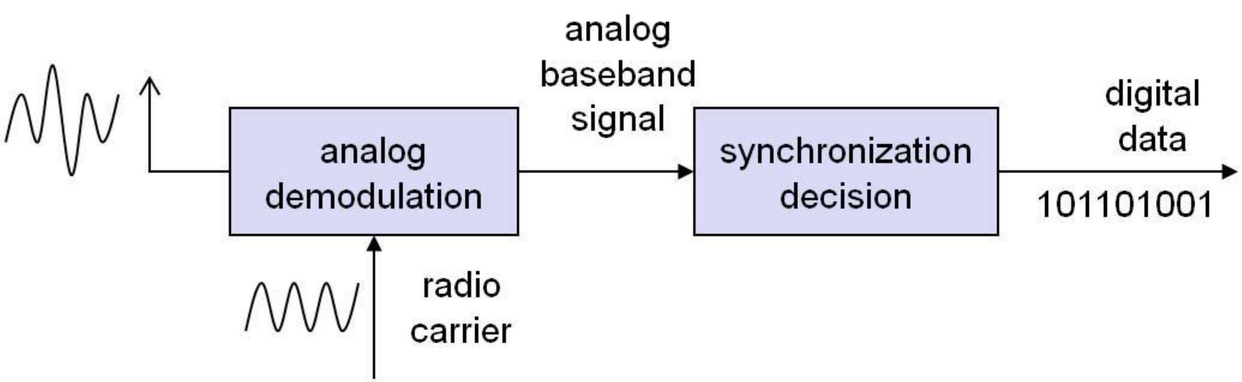

radio receiver

Modulation in action:

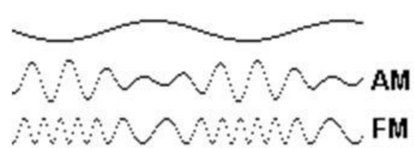

Digital modulation

- Modulation of digital signals known as Shift Keying

- Amplitude Shift Keying (ASK):
  - very simple
  - low bandwidth requirements
  - very susceptible to interference

- Frequency Shift Keying (FSK):
  - needs larger bandwidth

- Phase Shift Keying (PSK):
  - more complex
  - robust against interference

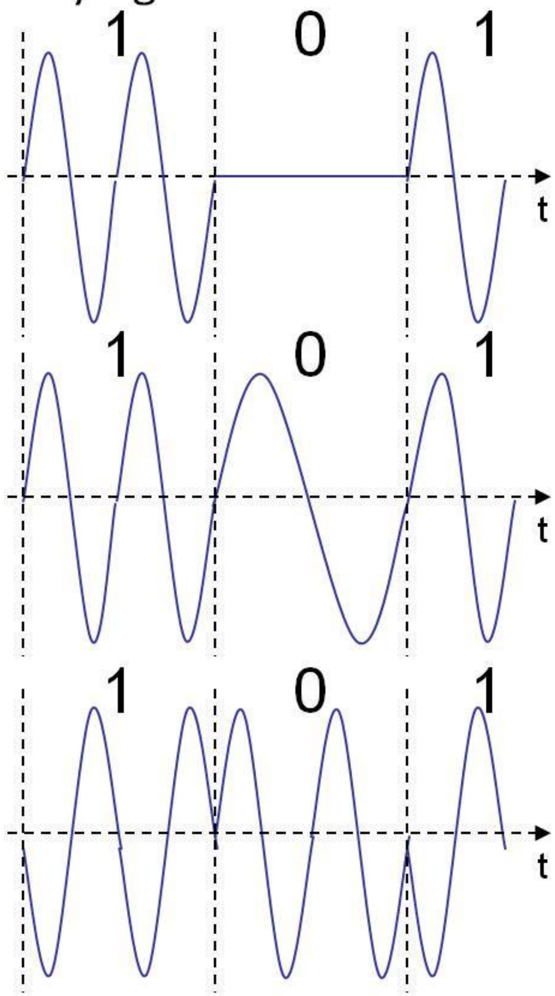

Phase Shift Keying 101

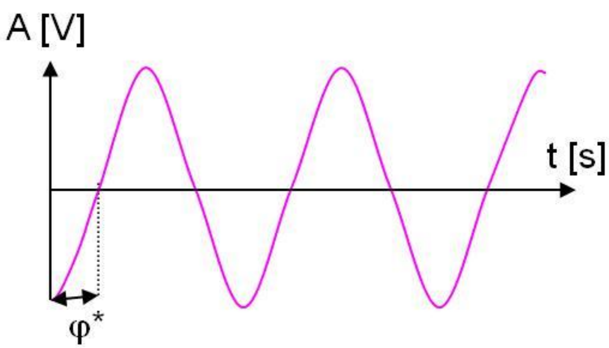

amplitude domain

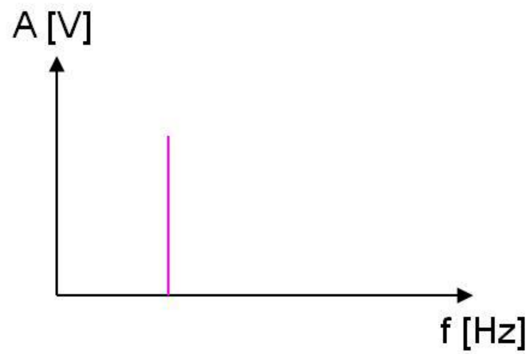

frequency spectrum

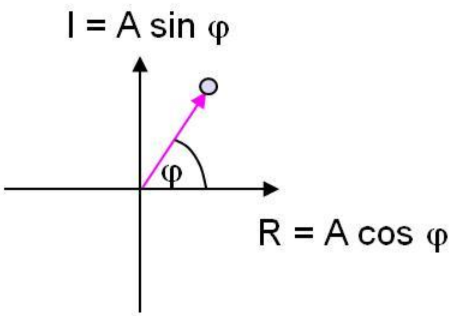

phase state diagram

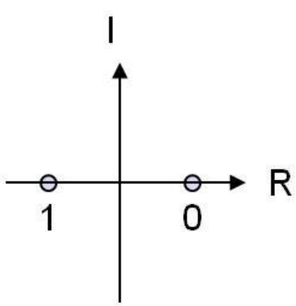

BPSK (robust, satellites)

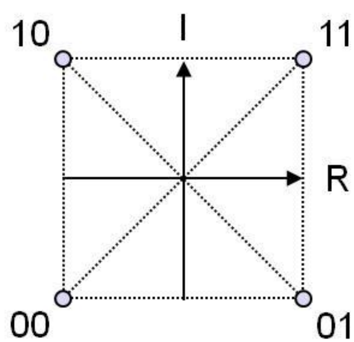

QPSK

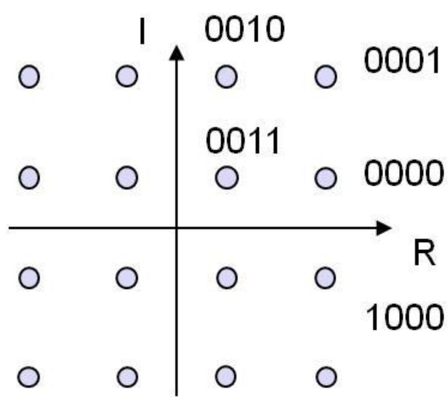

QAM (large β)

Signal-To-Interference-Plus-Noise Ratio (SINR) Formula

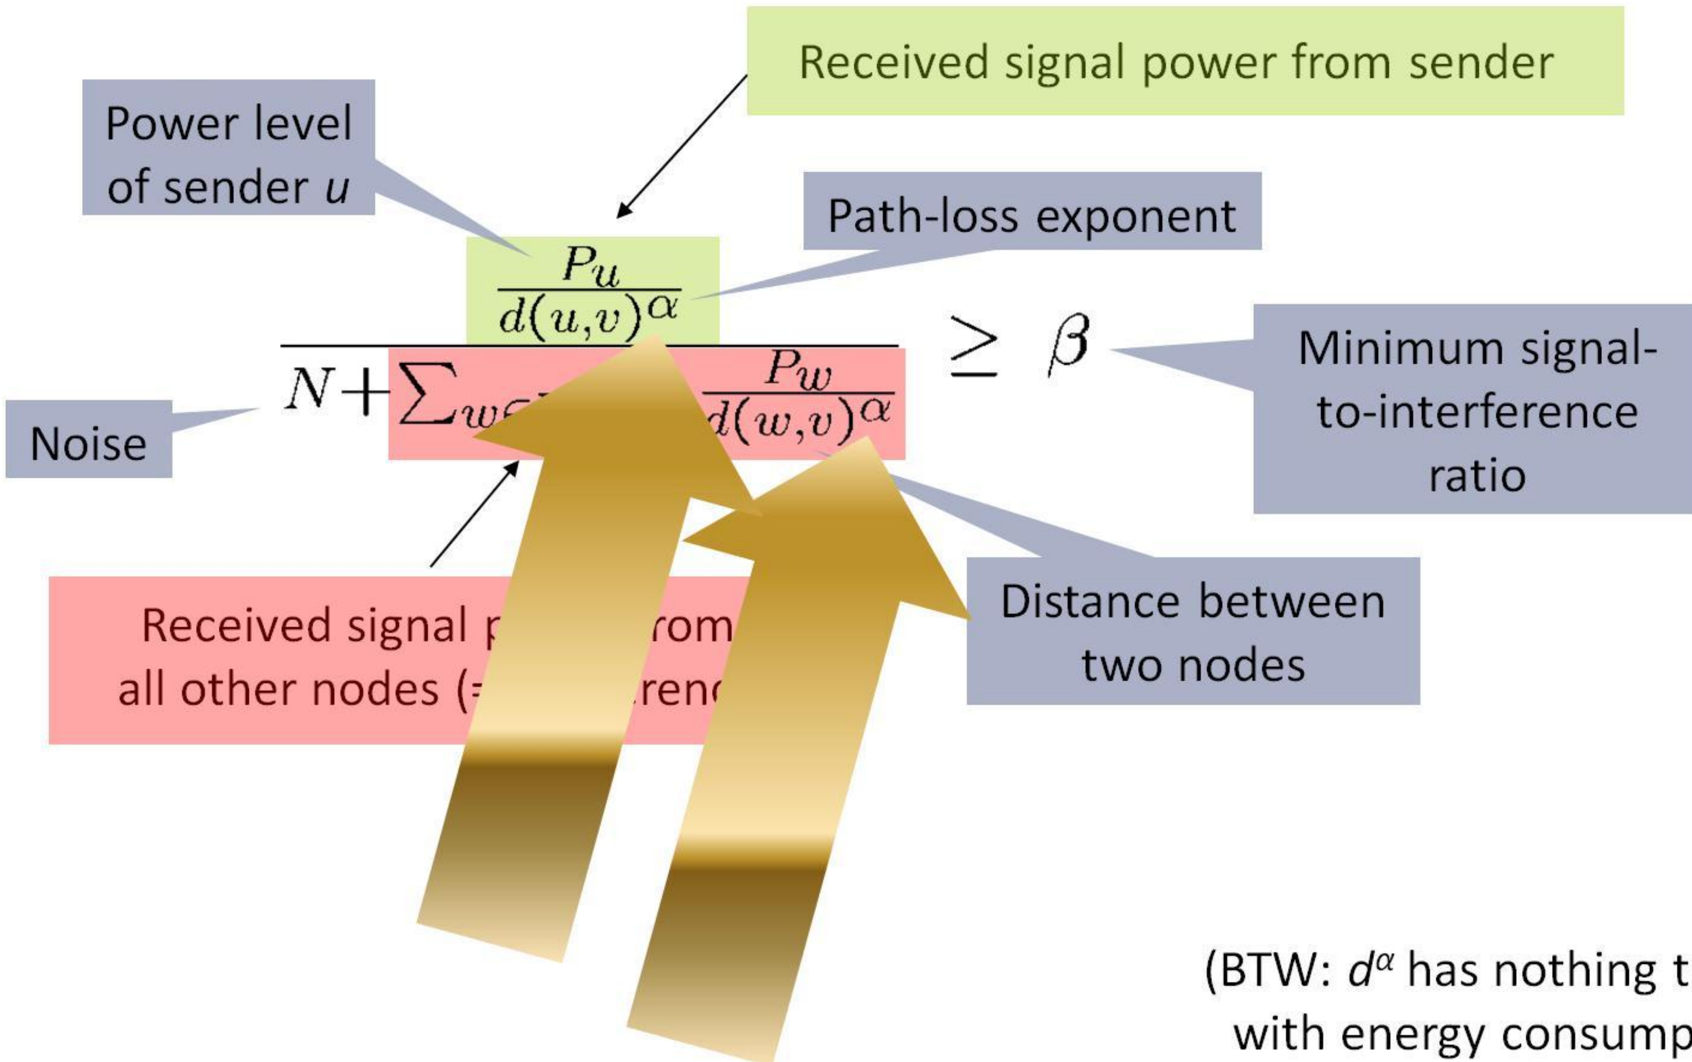

$$\frac{\frac{P_u}{d(u,v)^{\alpha}}}{N + \sum_{w \in \ldots} \frac{P_w}{d(w,v)^{\alpha}}} \geq \beta$$

(BTW: d^{α} has nothing to do with energy consumption)

Path-loss-exponent α

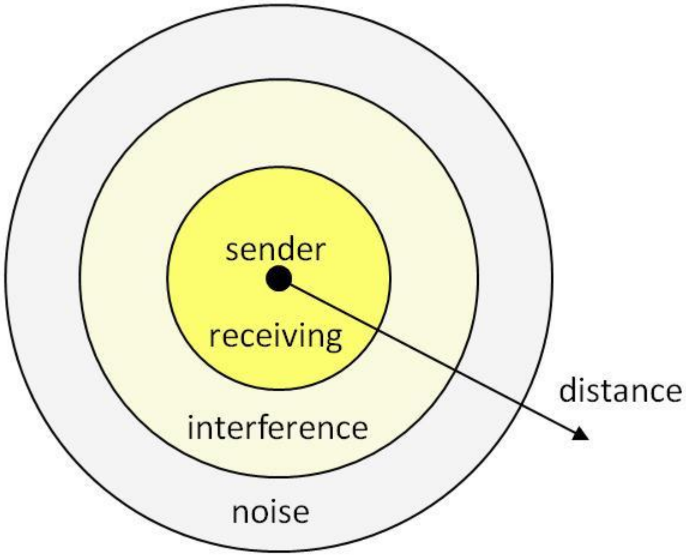

Path-loss-exponent α

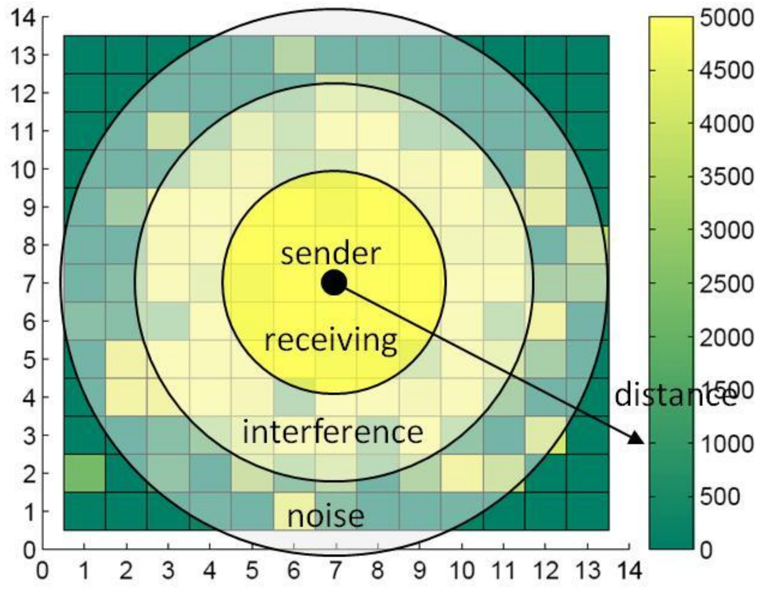

Path-loss-exponent α

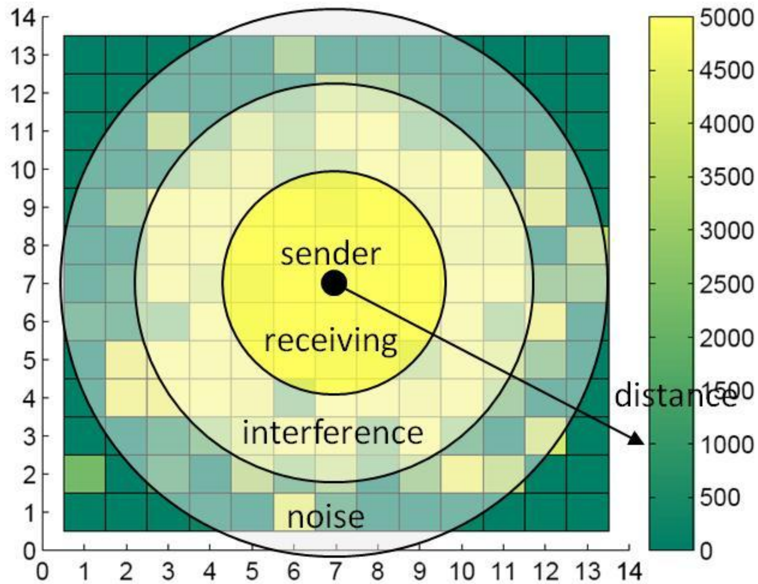

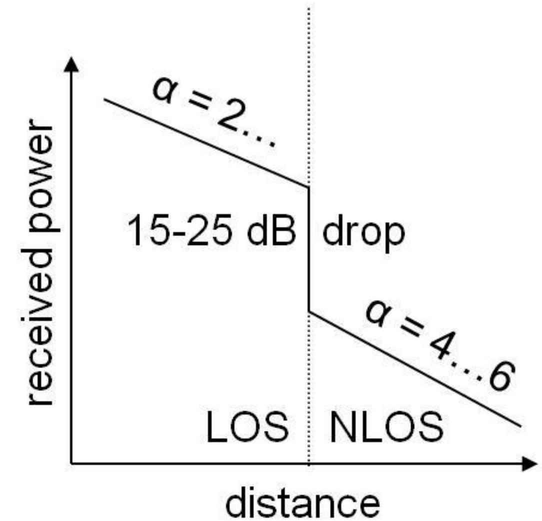

$$P_r = \frac{P_s G_s G_r \lambda^2}{(4\pi)^2 d^2 L}$$

$$P_r = \frac{P_s G_s G_r h_s^2 h_r^2}{d^4}$$

Path-loss-exponent α

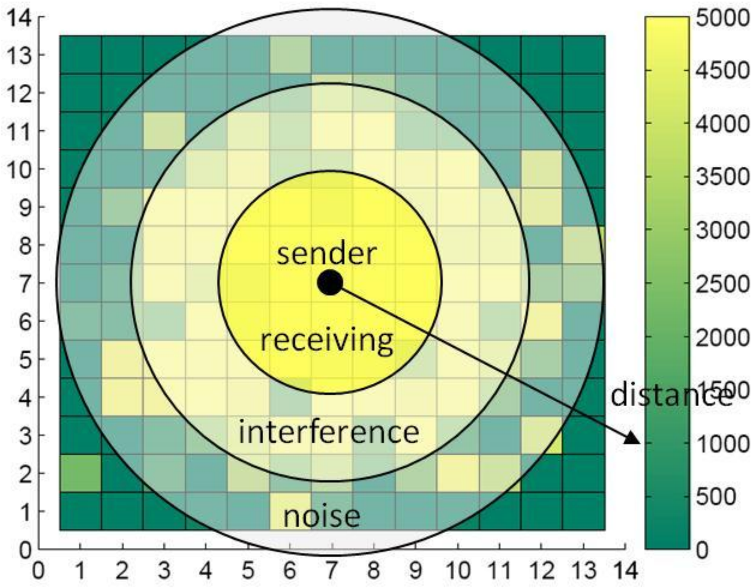

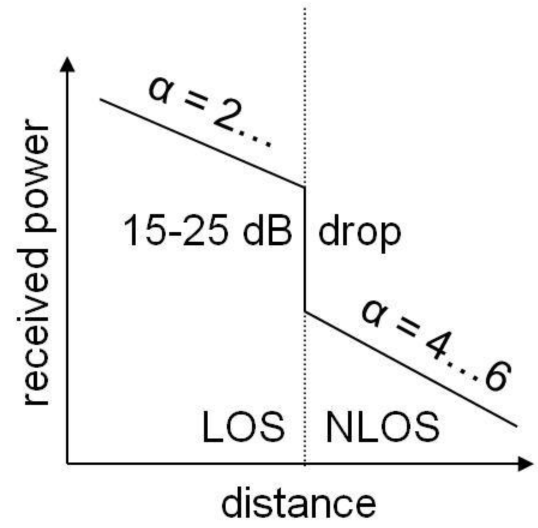

$$P_r = \frac{P_s G_s G_r \lambda^2}{(4\pi)^2 d^2 L}$$

$$P_r = \frac{P_s G_s G_r h_s^2 h_r^2}{d^4}$$

$\alpha \geq$ Dimension

2nd law of thermodyn.

Wireless Propagation Depends on Frequency

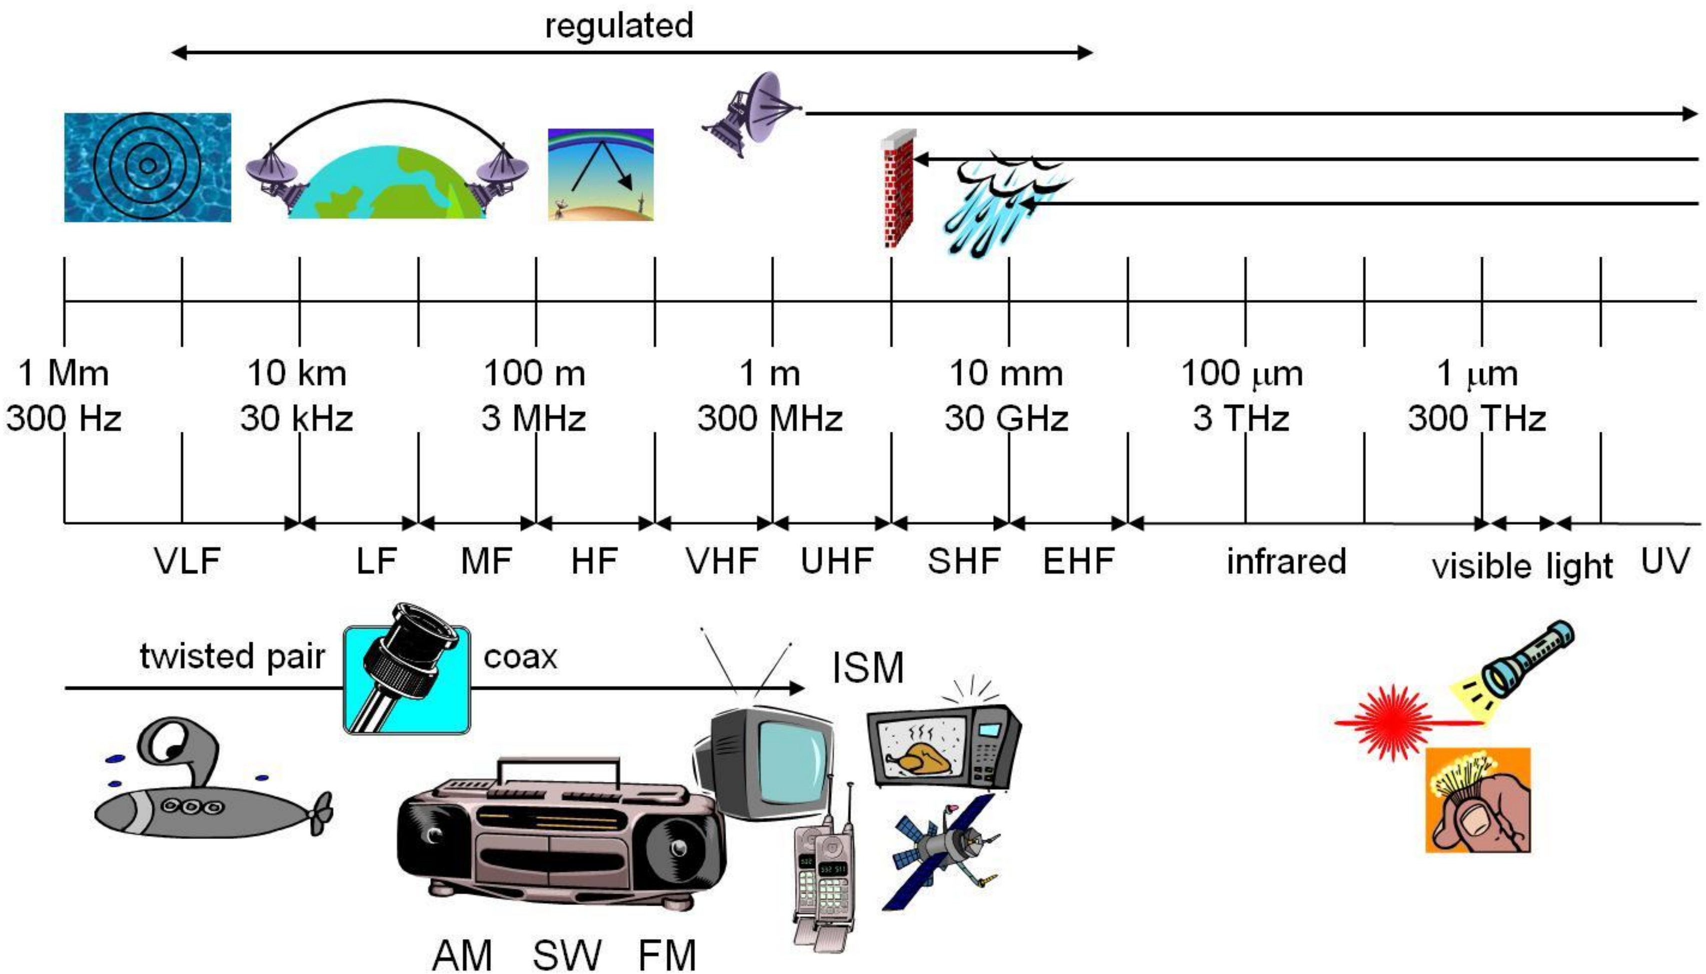

Path-loss-exponent α: Near-Field Effects

$$P_r = \frac{P_s G_s G_r \lambda^2}{(4\pi)^2 d^2 L}$$

d << 1?

Path-loss-exponent α: Near-Field Effects

$$P_r = \frac{P_s G_s G_r \lambda^2}{(4\pi)^2 d^2 L}$$

d << 1?

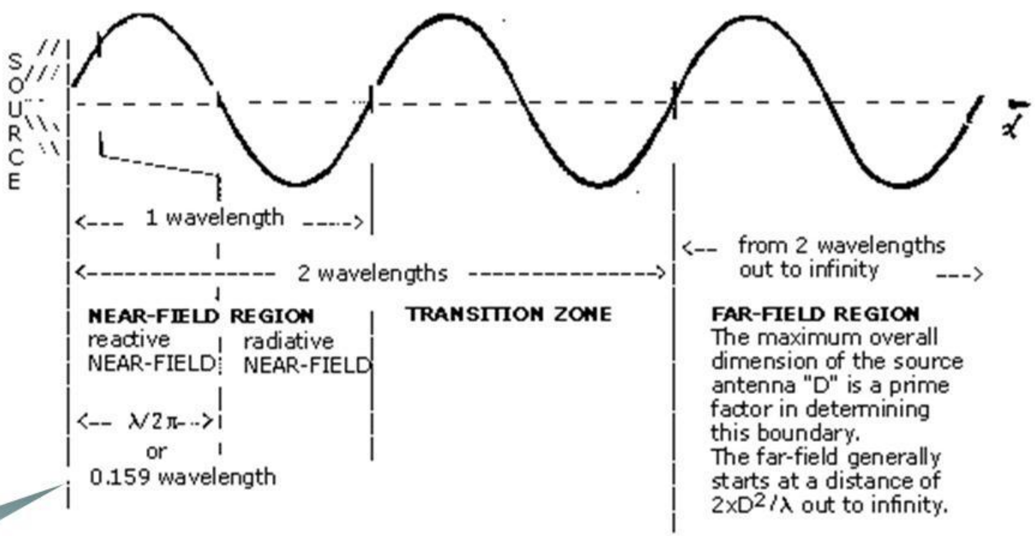

1cm!!

… in other words, algorithmic papers should rule out near-field effects

Real World Examples

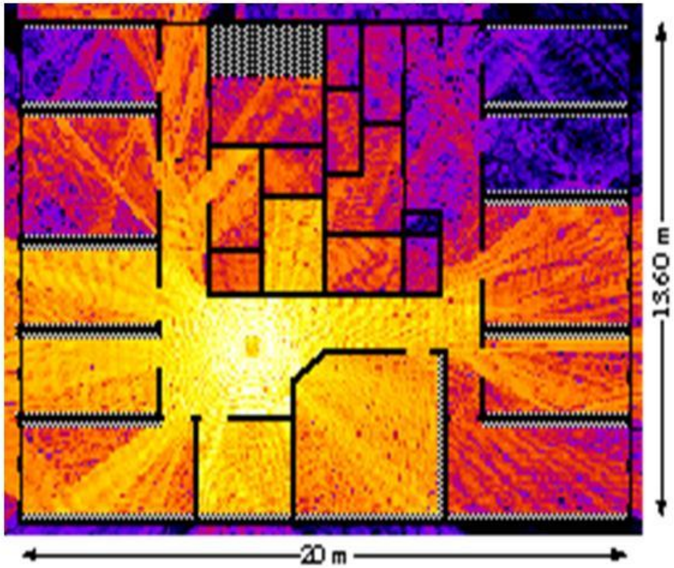

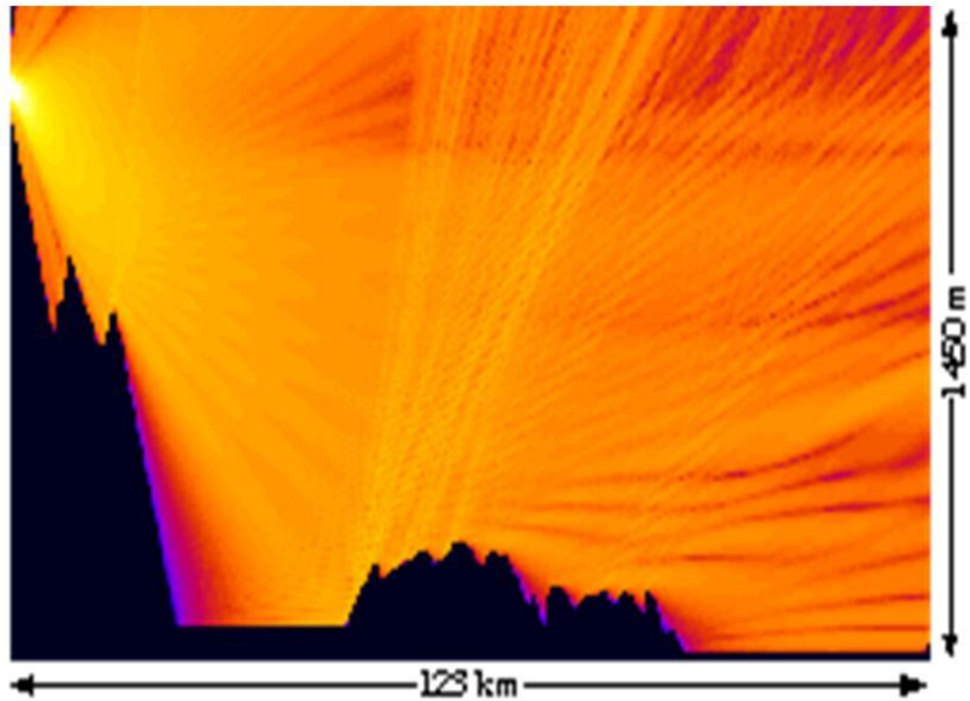

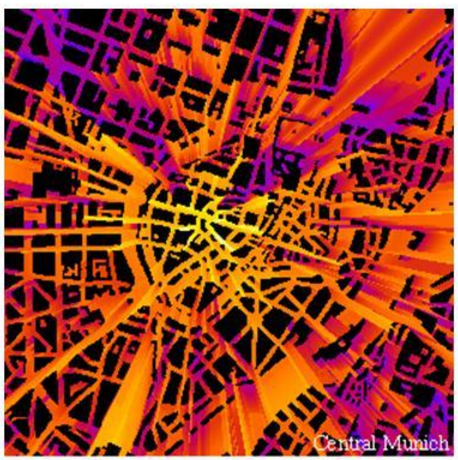

Attenuation by objects

- Shadowing (3-30 dB):
  - textile (3 dB)
  - concrete walls (13-20 dB)
  - floors (20-30 dB)
- reflection at large obstacles
- scattering at small obstacles
- diffraction at edges
- fading (frequency dependent)

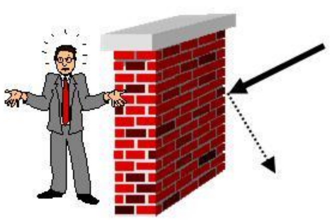

shadowing

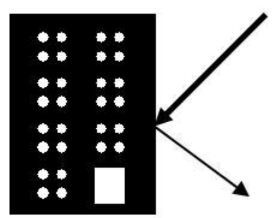

reflection

scattering

diffraction

Multipath

- Signal can take many different paths between sender and receiver due to reflection, scattering, diffraction

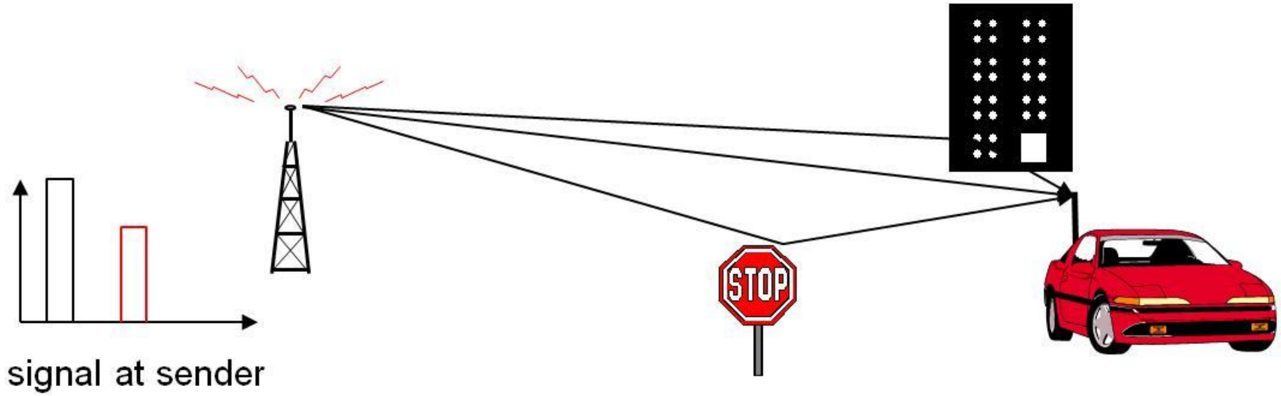

- Time dispersion: signal is dispersed over time

Multipath

- Signal can take many different paths between sender and receiver due to reflection, scattering, diffraction

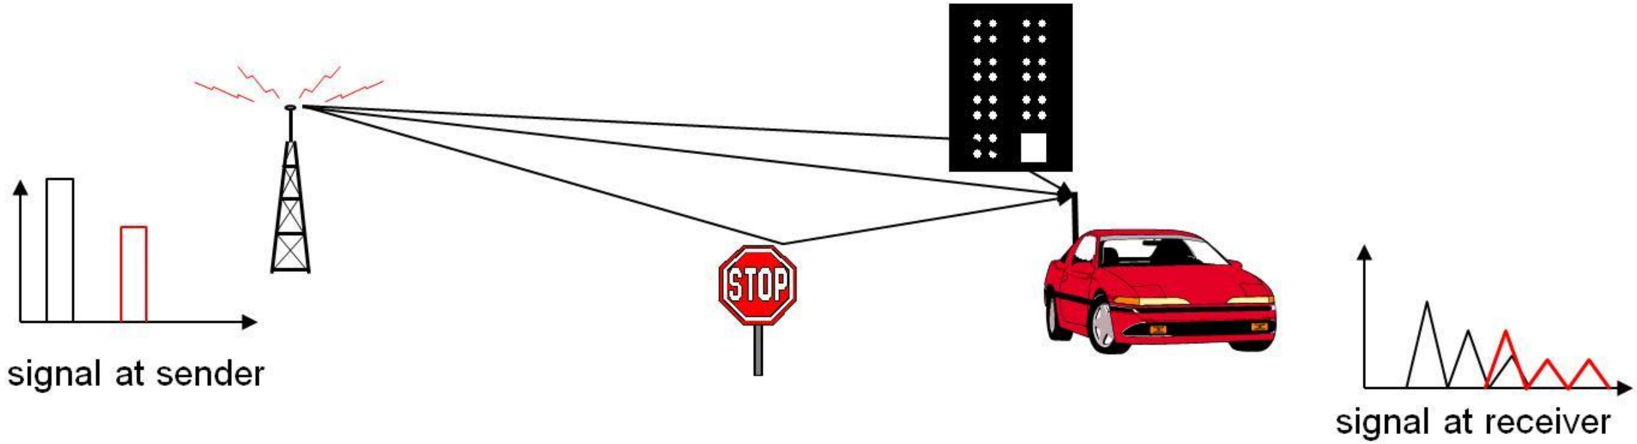

- Time dispersion: signal is dispersed over time
- Interference with “neighbor” symbols: Inter Symbol Interference (ISI)

Multipath

- Signal can take many different paths between sender and receiver due to reflection, scattering, diffraction

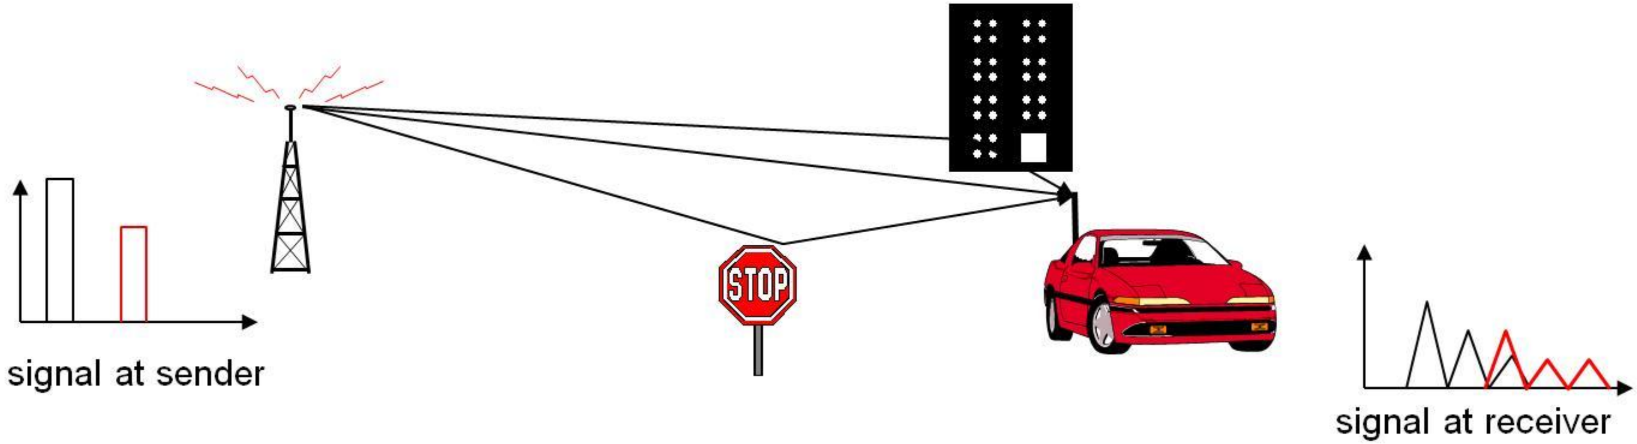

- Time dispersion: signal is dispersed over time
- Interference with “neighbor” symbols: Inter Symbol Interference (ISI)
- The signal reaches a receiver directly and phase shifted. Distorted signal depending on the phases of the different parts

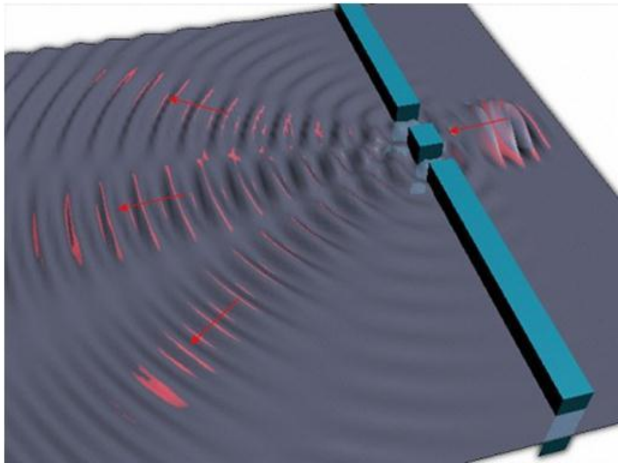

There’s much more…

MIMO

There's much more…

MIMO

Analog Network Coding

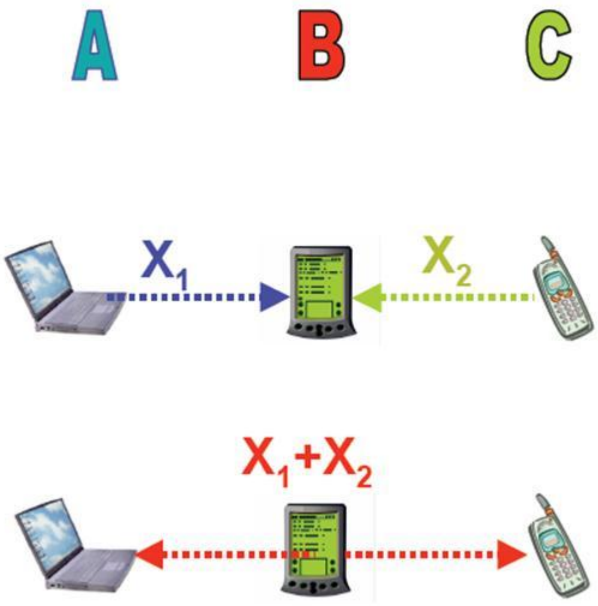

Advanced Algorithms?

Quiz: How to Build a Multi-Hop Alarm System

Problem: More than 1 node may sense problem at the same time.

Potentially we have a massive interference problem!

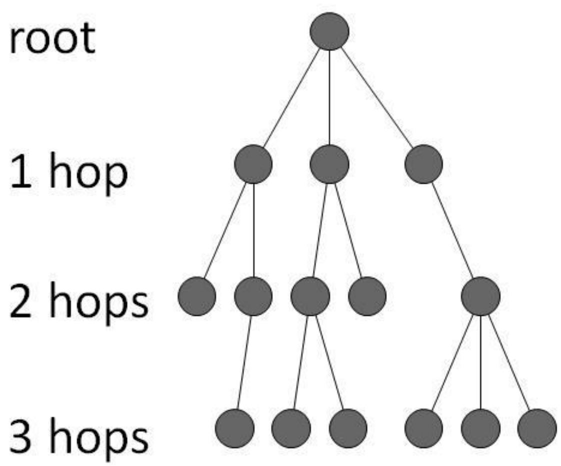

Quiz: How to Build a Multi-Hop Alarm System

Problem: More than 1 node may sense problem at the same time.

Potentially we have a massive interference problem!

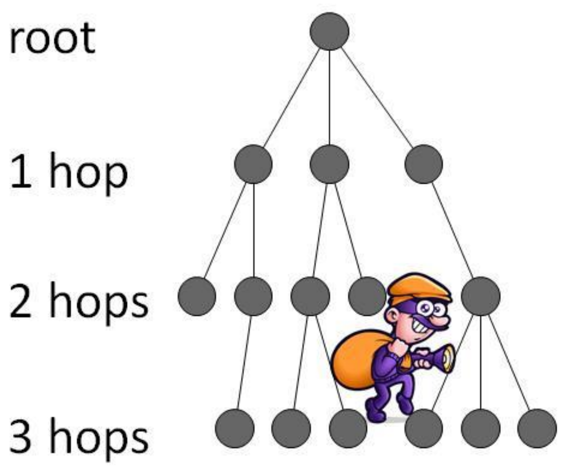

Quiz: How to Build a Multi-Hop Alarm System

Problem: More than 1 node may sense problem at the same time.

Potentially we have a massive interference problem!

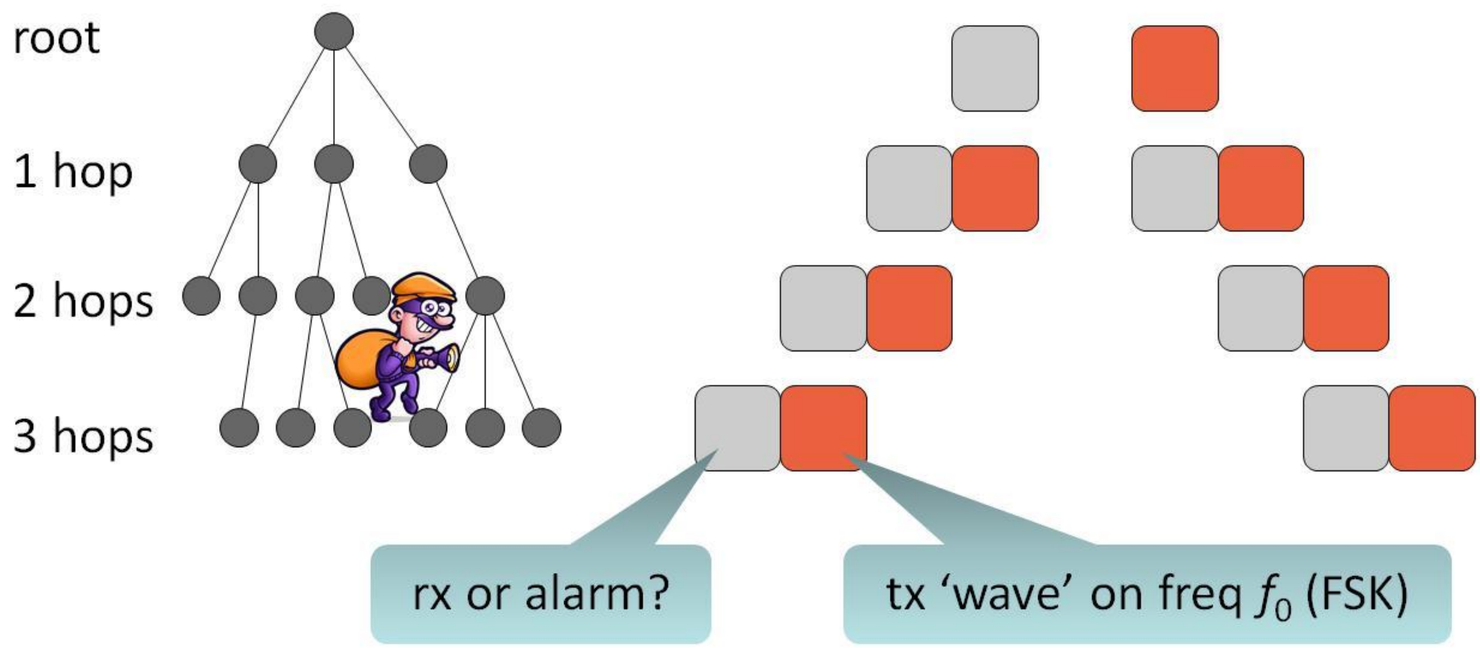

…followed by verification (1% false positives outdoors)

[Flury, W, 2010]

Media Access: Theory and Practice

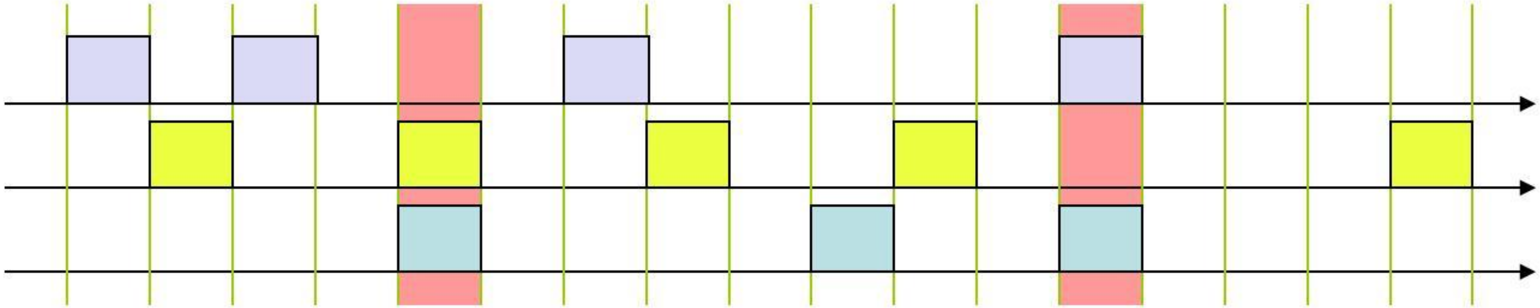

[Sommer, W, 2010]

Media Access: Theory and Practice

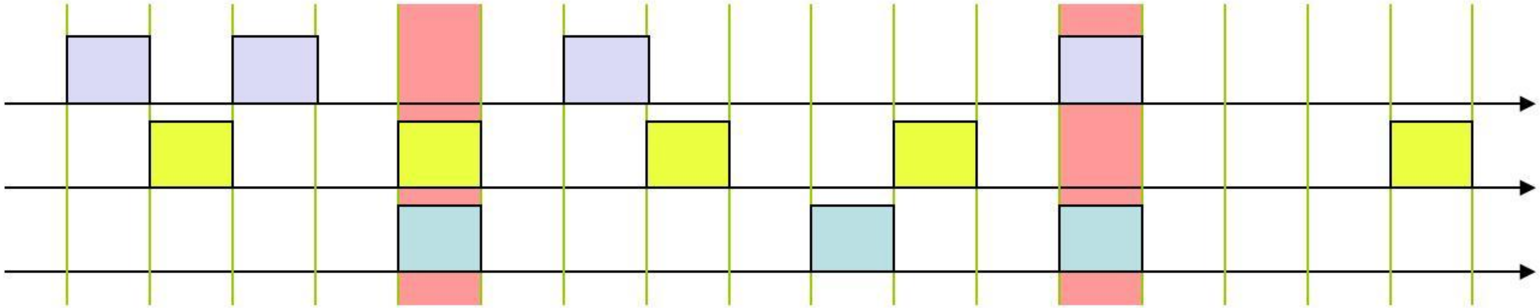

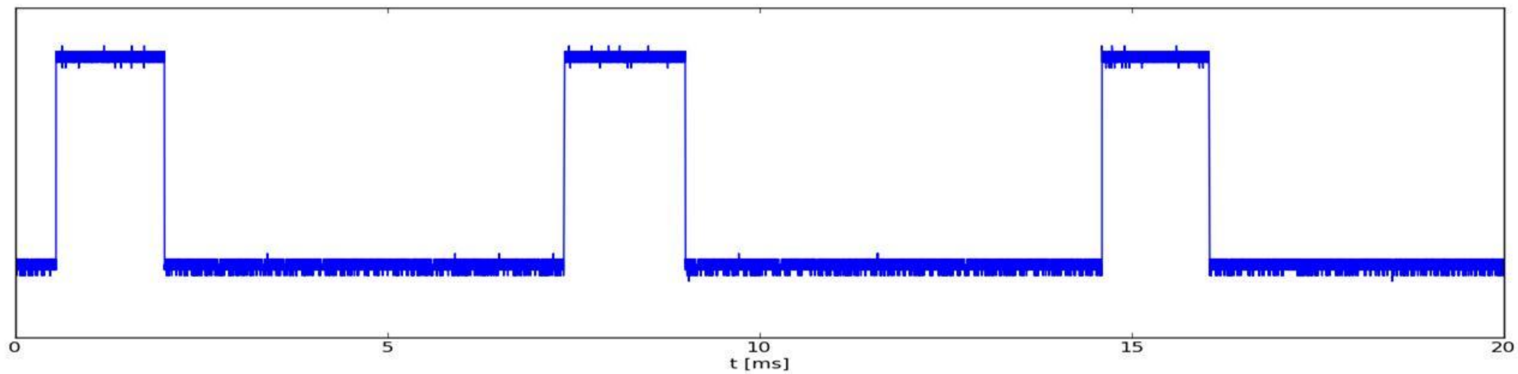

[Sommer, W, 2010]

Media Access: Theory and Practice

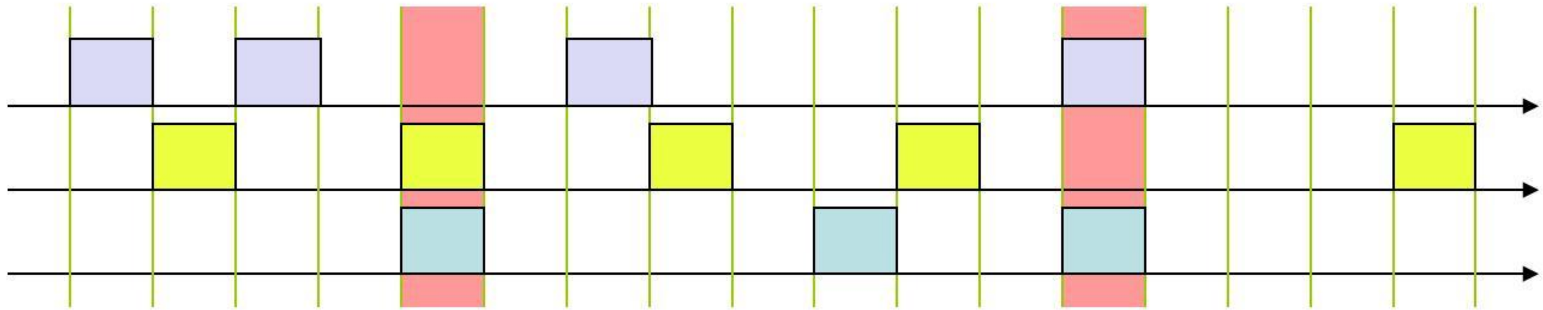

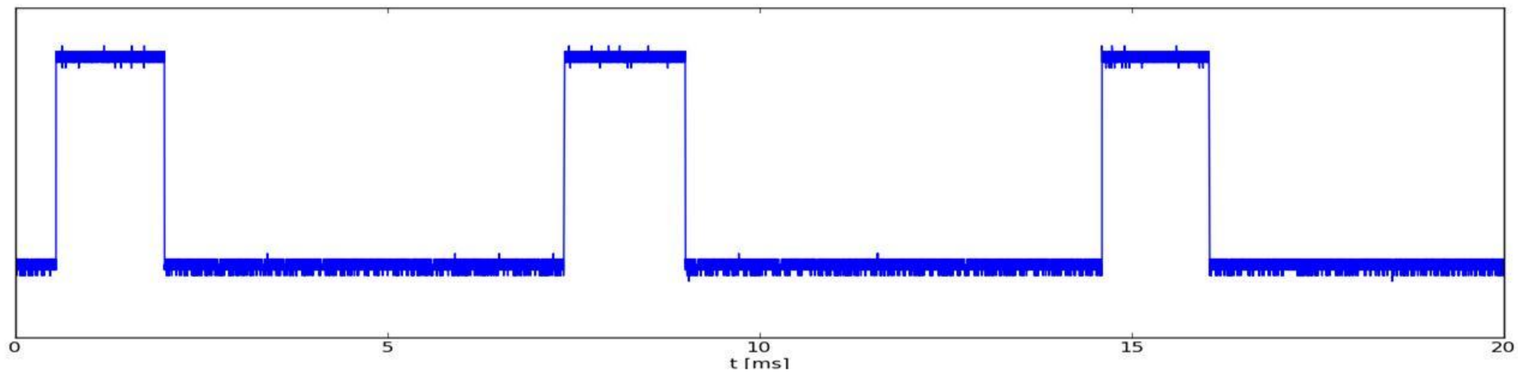

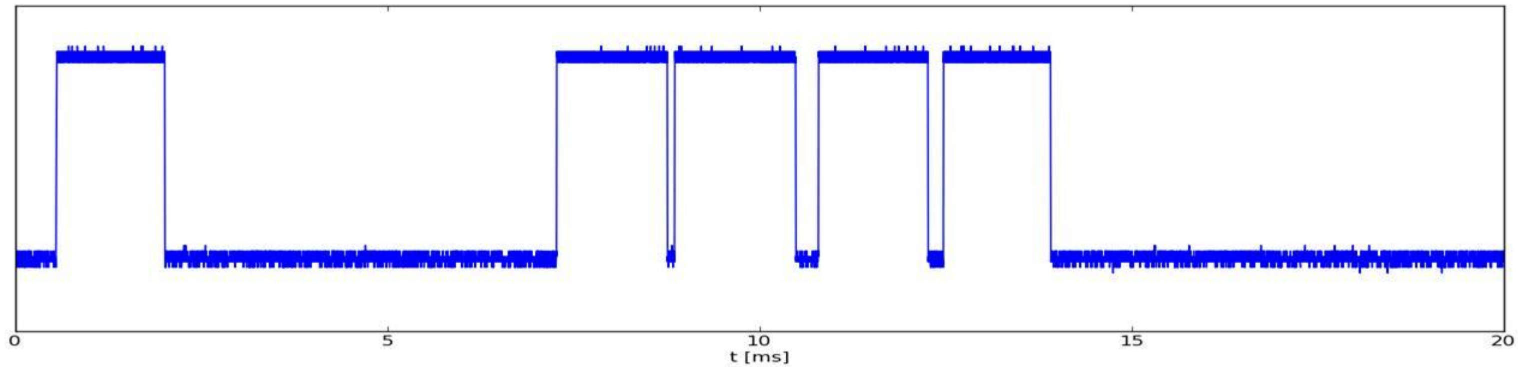

[Sommer, W, 2010]

Ultrasound

(A different kind of communication)

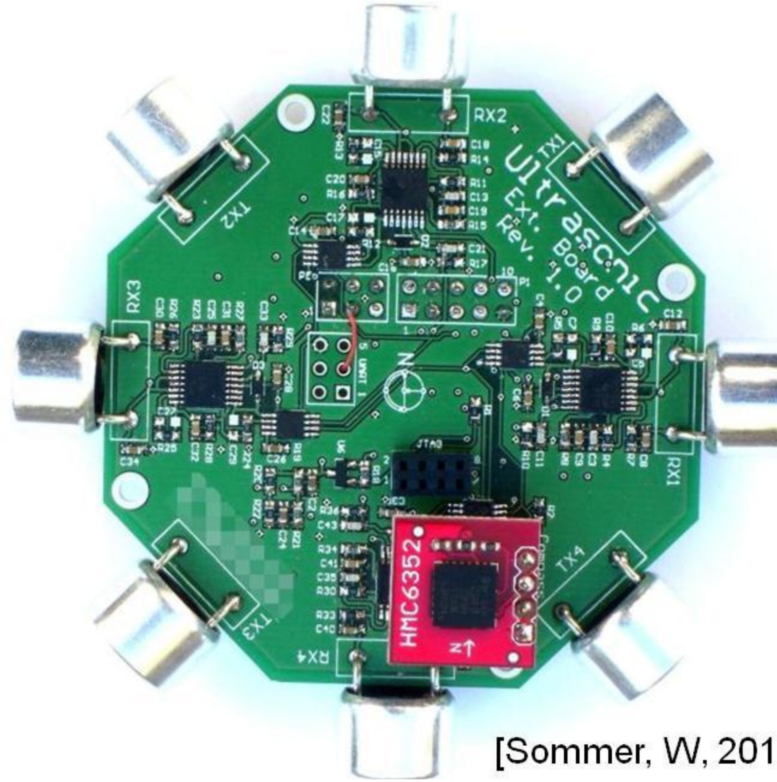

[Sommer, W, 2010]

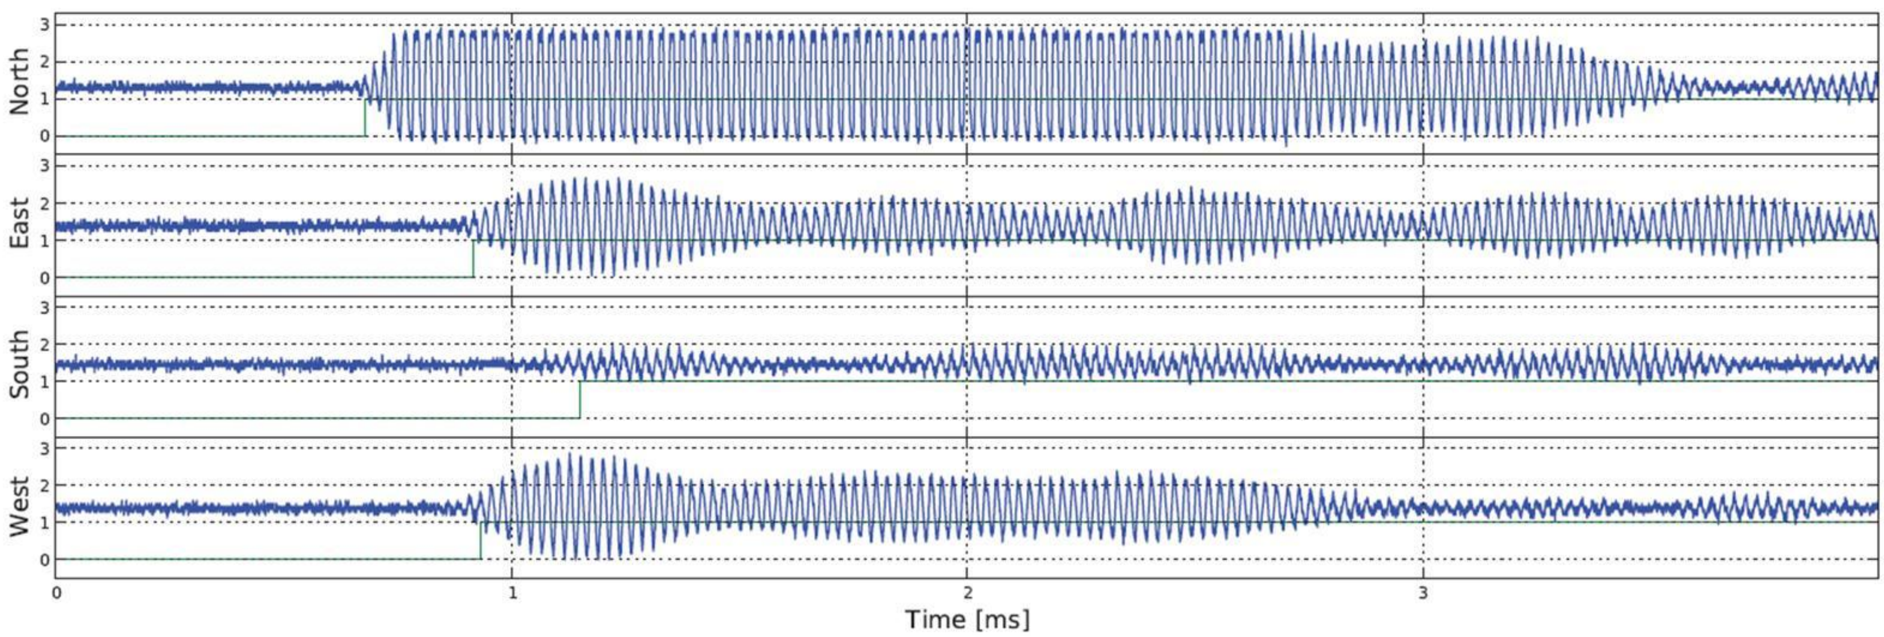

Ultra-Wideband (UWB)

- An example of a new physical paradigm.
- Discard the usual dedicated frequency band paradigm.
- Instead share a large spectrum (about 1-10 GHz).

- Modulation: Often pulse-based systems. Use extremely short duration pulses (sub-nanosecond) instead of continuous waves to transmit information. Depending on application 1M-2G pulses/second

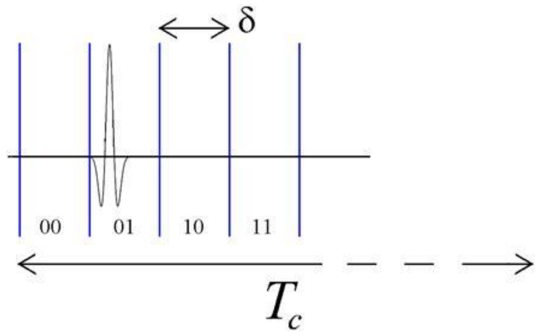

PPM

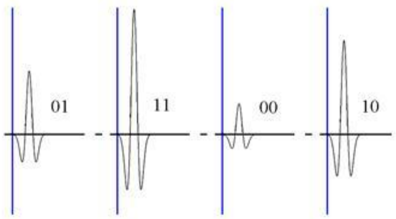

PAM

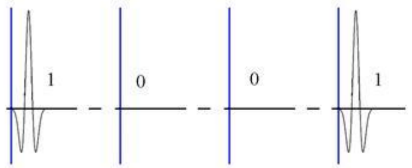

OOK

Summary

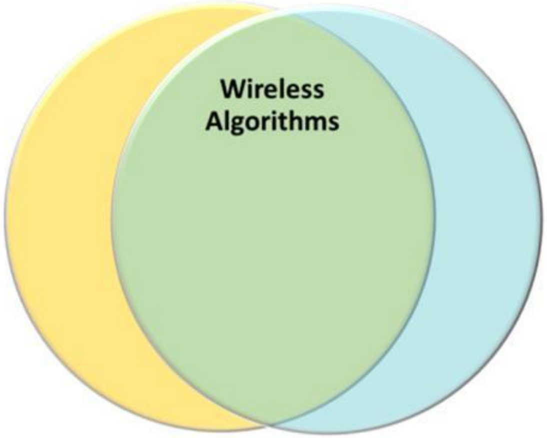

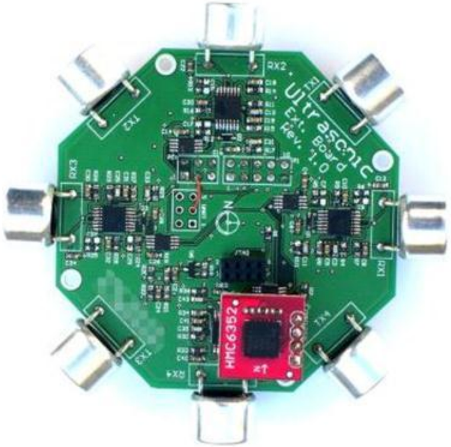

Thank You!

Questions & Comments?

Roger Wattenhofer