RUTGERS

THE STATE UNIVERSITY OF NEW JERSEY

The Case for Transmit Only Communication

» Presented by Rich Martin

» And many more, including
» Richard Howard, Yanyong Zhang,
» Giovanni Vannuci, Junichiro Fukuyama,
» Bernhard Firner, Robert Moore, Chenren Xu

»

The Opportunity

Moore's Law: # transistors on cost-effective chip doubles every 18 months

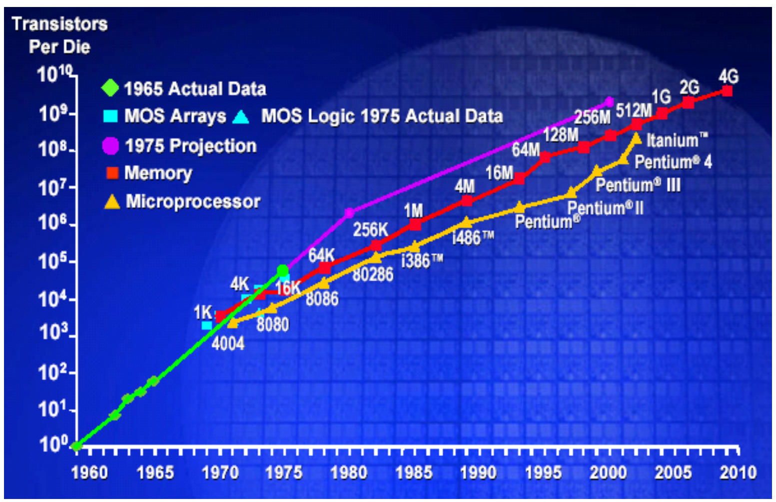

Today: 1 million transistors per $

Bell's Law: a new computer class emerges every 10 years

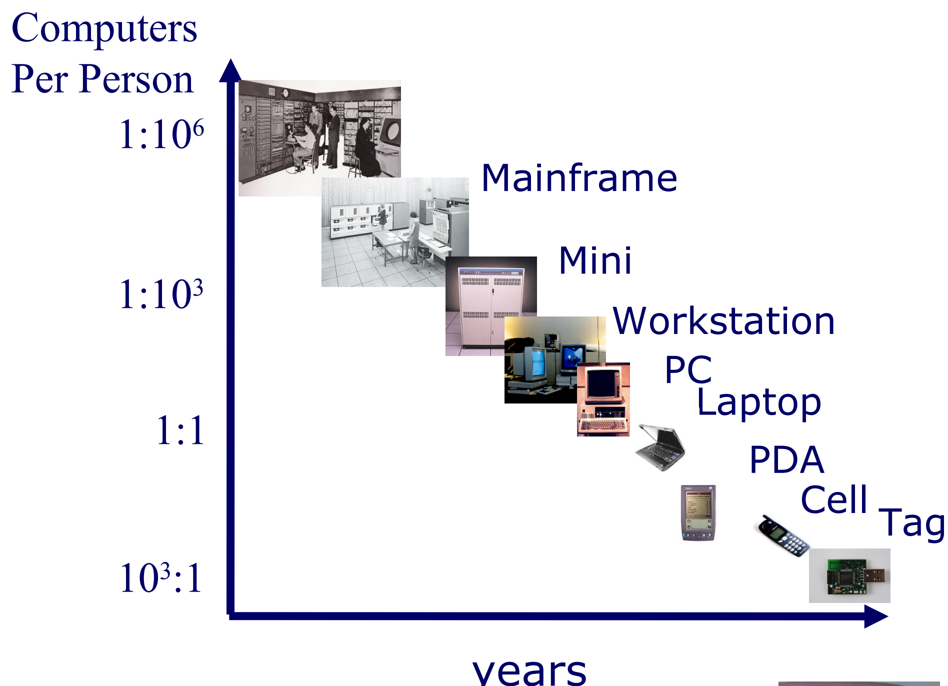

Same fabrication technology provides CMOS radios for communication and micro-sensors

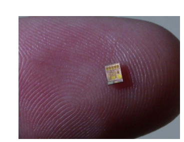

The Vision

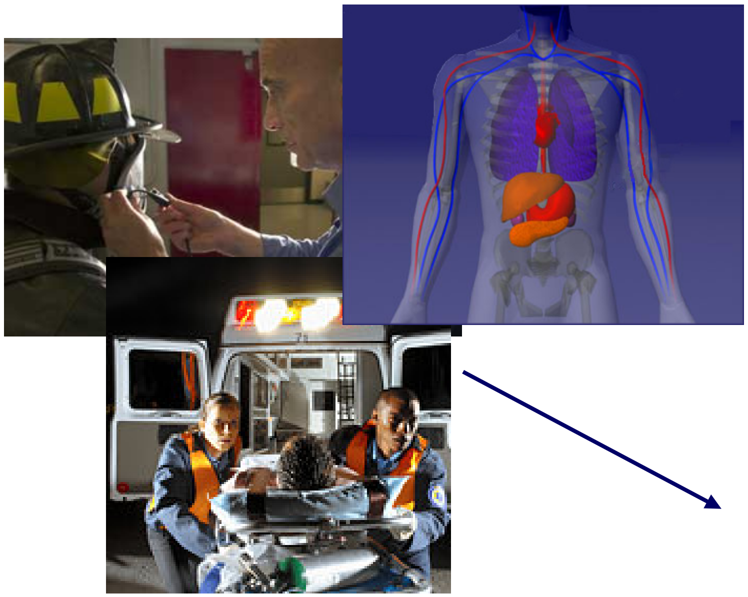


Name = 16
Parent = 0
Epoch = 244
Time = 2003-09-22 21:09:12
What's happening?
Timeline
@Mentions
Retweets
Searches
Lists
summerfrog2 summerfrog2
1 cup poured --message at Fri Mar 25 10:59:13 EDT 2011
25 Mar
summerfrog2 summerfrog2
Coffee is ready! --message at Fri Mar 25 10:52:46 EDT 2011
25 Mar Favorite Reply Delete
summerfrog2 summerfrog2
3 cups poured --message at Fri Mar 25 10:51:18 EDT 2011
25 Mar
summerfrog2 summerfrog2
2 cups poured --message at Fri Mar 25 10:50:07 EDT 2011
25 Mar
summerfrog2 summerfrog2
1 cup poured --message at Fri Mar 25 10:47:26 EDT 2011
25 Mar
summerfrog2 summerfrog2
25 Mar
close
@summerfrog2
summerfrog2
Coffee is ready! --message at Fri Mar 25 10:52:46 EDT 2011
25 Mar via LabTesting Favorite Reply Delete
Tweets from @summerfrog2
summerfrog2 summerfrog2
1 cup poured --message at Fri Mar 25 10:59:13 EDT 2011
25 Mar

RUTGERS

Down the garden path of sensor networks

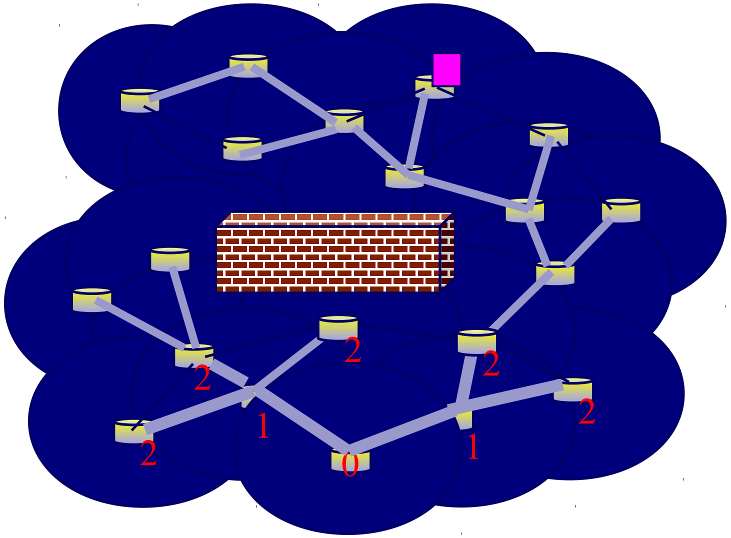

Programming a sensor network:

- Multi-hop
- Ad-hoc
- Aggregation and compression
- Energy conservation of whole application is paramount
- Novel operating systems, programming languages and environments

RUTGERS

A rose by any other name

- 1999 Smart Dust
- 2000 Sensor Networks
- 2004 Internet of Things
- 2005 Ambient Intelligence
- 2009 Swarms

- ~15 years on, we still have not realized the vision.

What happened?

Problems

- Problems people talked about:
  - Energy conservation
  - Scaling number of sensors
  - Efficiency of code data size in small sensors
  - Routing
- More meaningful problems:
  - Too expensive for application domains
  - Difficult to develop applications
  - Can't re-use infrastructure
  - Not general purpose

When less is better

- 1 level of the system performs 1 goal
  - Move other functionality to other layers
  - Overall system improvement!
- Architecture: RISC vs. CISC
  - Focus on instruction throughput, move abstraction to language/compiler
- Networks: IP vs. ISDN/ATM
  - Focus on packet switching, move circuits and sessions to endpoints
- Operating systems: Unix vs. Multics
  - Focus on process execution and I/O, move object persistence to the database

Transmit Only Approach

- Key insight: sensed data is in a class where small losses can be tolerated. Probabilistic reception is OK.
  - Similar to audio, video, and multi-player games, not documents.
- Sensors only sense and transmit with specified periods
  - Sensors are at most 1 hop
  - Add small amount of randomization to prevent collision periodicity.
- A small set of receivers cooperate to reconstruct sensed data
  - Connected by a powerful back-haul network
  - Back-haul bandwidth > sensor bandwidth

TO - less is better

- Everything that doesn't transmit an application bit is overhead
- Removed:
  - Sensing the channel before transmission (for CSMA protocols)
  - Acknowledgements (for RTS/CTS protocols)
  - Precise clocks and synchronization (for TDMA protocols)
  - Signal feedback (MIMO physical layers)

Transmit Only as less is better

- Focus on getting the sensed data:
  - Everything else is overhead
- Saves energy on the sensor
  - Receiving has similar energy costs per bit-time as transmit
- Simplify the sensors
  - Fewer components
  - Cheaper components
  - Smaller sensors
- Simply the programming model
  - Aggregation layer's interface to the sensors becomes much simpler.

Challenges

- Semantic: Some data loss OK?
- Wireless Channel Utilization
  - Are we really limited to 18% efficiency?
- Receiver Network:
  - Complexity?
  - Energy use?
  - Number and coverage?
  - Connectivity?
- Manageability
  - Change parameters on the sensor?
- Security
  - How to perform lightweight unidirectional security?

Outline

- Improvements from sensor simplification
- Recovering channel efficiency exploiting the capture effect
- Simplifying the programming model
- Example applications

Sensor simplicity

- Sensor node cost is a limitation for many applications
  - Applications enabled at sensor cost of $100, $10, $1, 10¢, 1¢ ?
- Cost assumptions based on scaling Moore's law real omit real constraints
  - 15 years show these constraints are fundamental
- Cost is driven by the number and type of components, not Moore's law!
- TO reduces costs by several factors
  - enough to expand the application space ($80->$10)
- Marginal costs will only go down if there is a true single-chip sensor
  - But high fixed costs remain a barrier for a true single chip solution!

Two wireless sensor boards

Classic

TelosB (2004)

Transmit-Only

TO-PIP(2013)

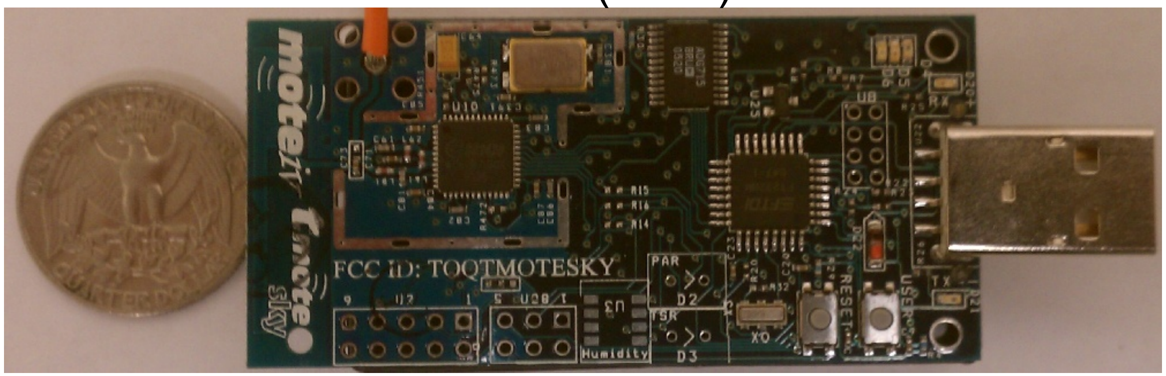

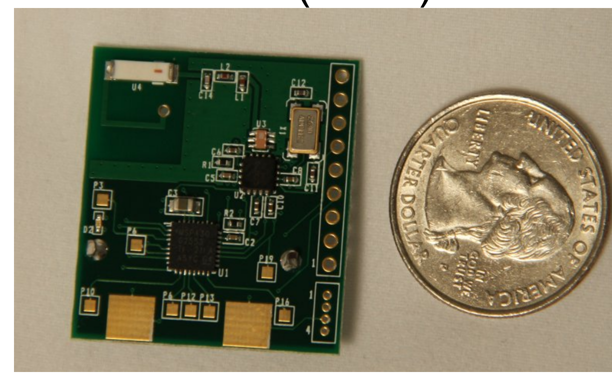

Antenna

Radio

Micro controller

Battery

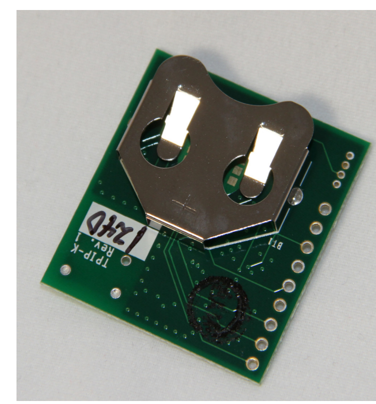

Component counts

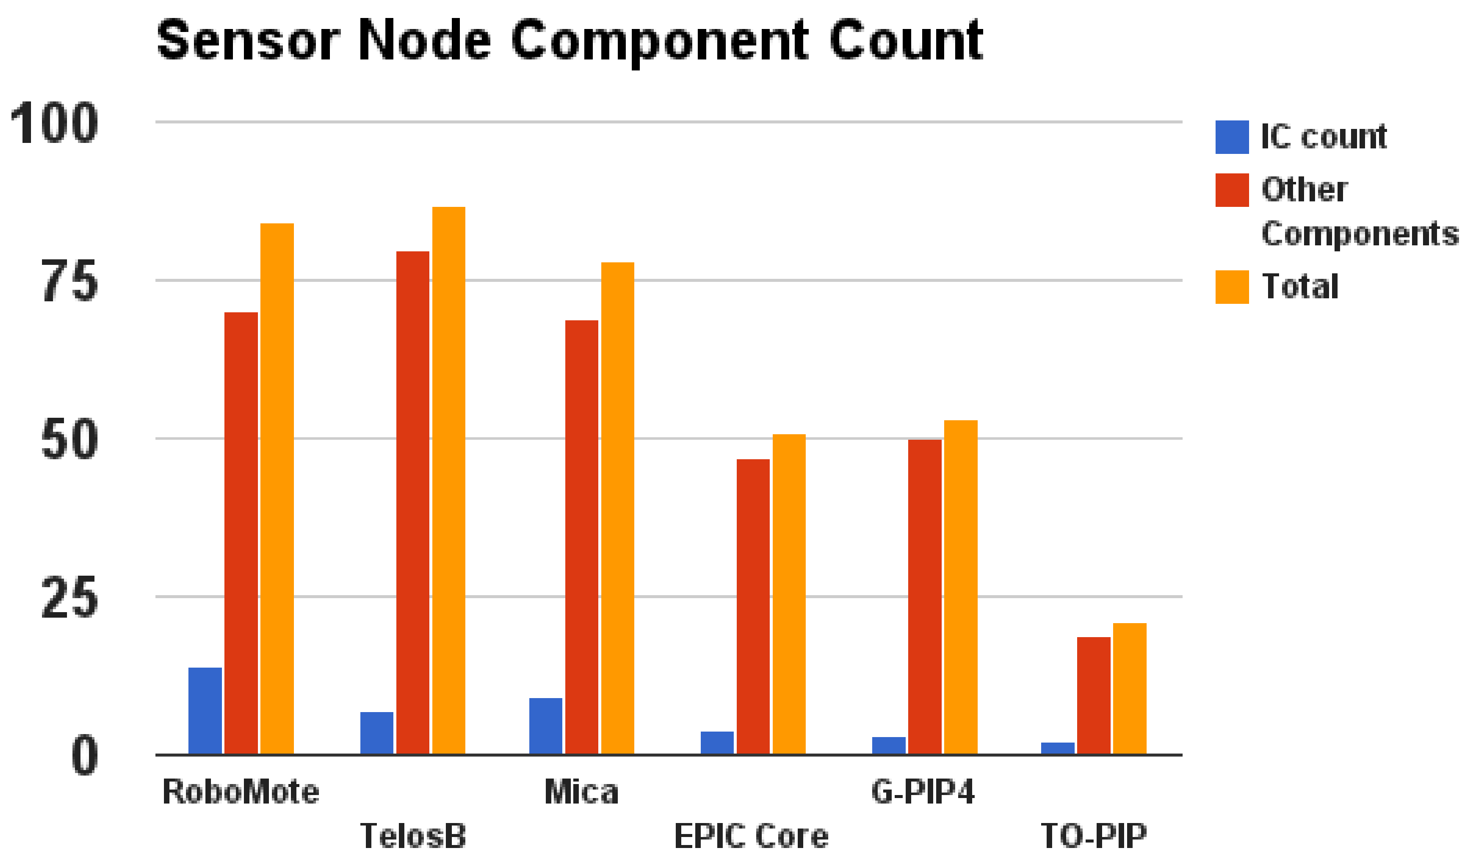

RUTGERS

Component Cost vs. Volume

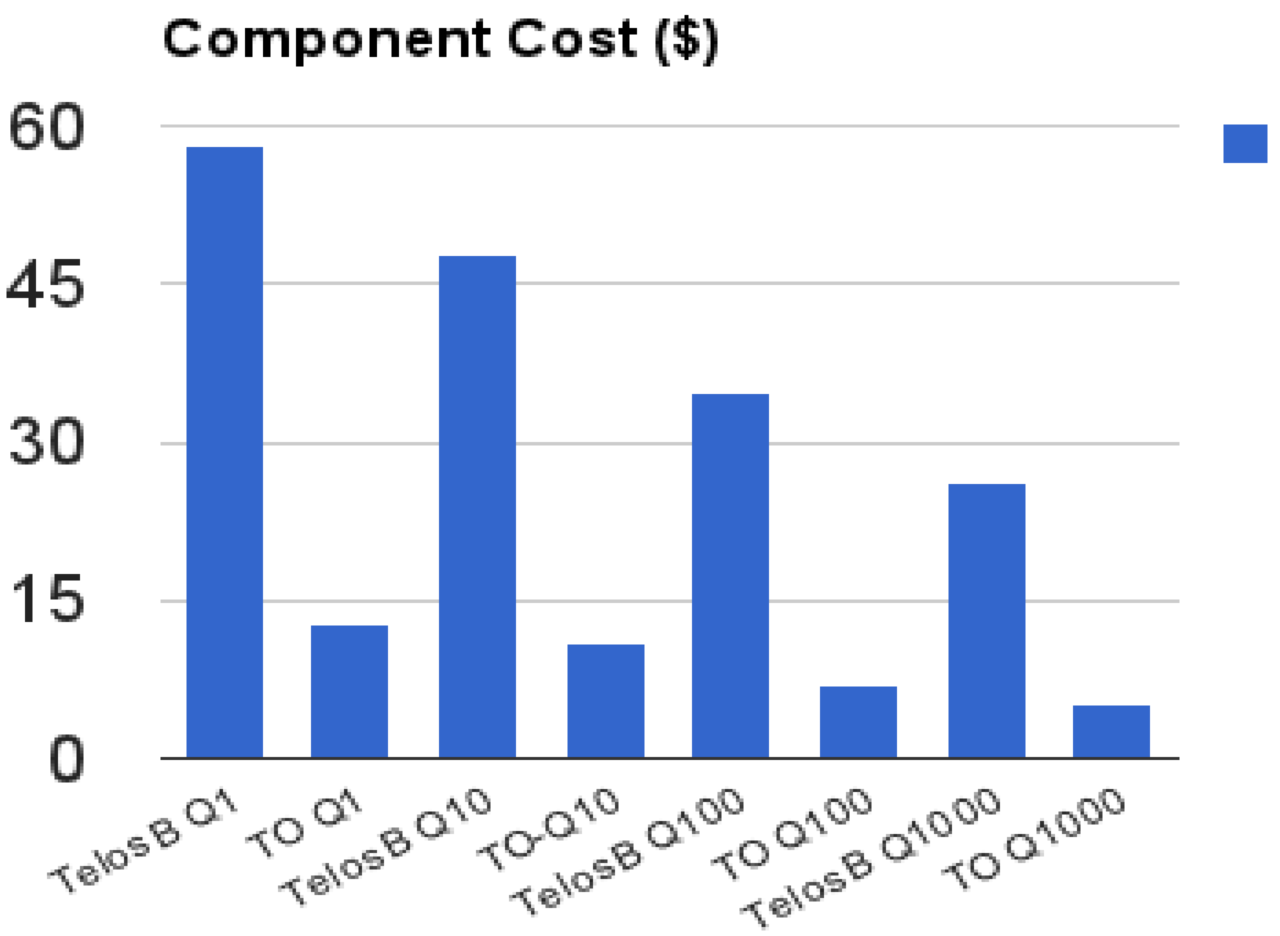

Problems with existing Systems on a Chip

- CPU+Radio: 1 chip, but 39 components
- Analog components do not scale with Moore's Law!

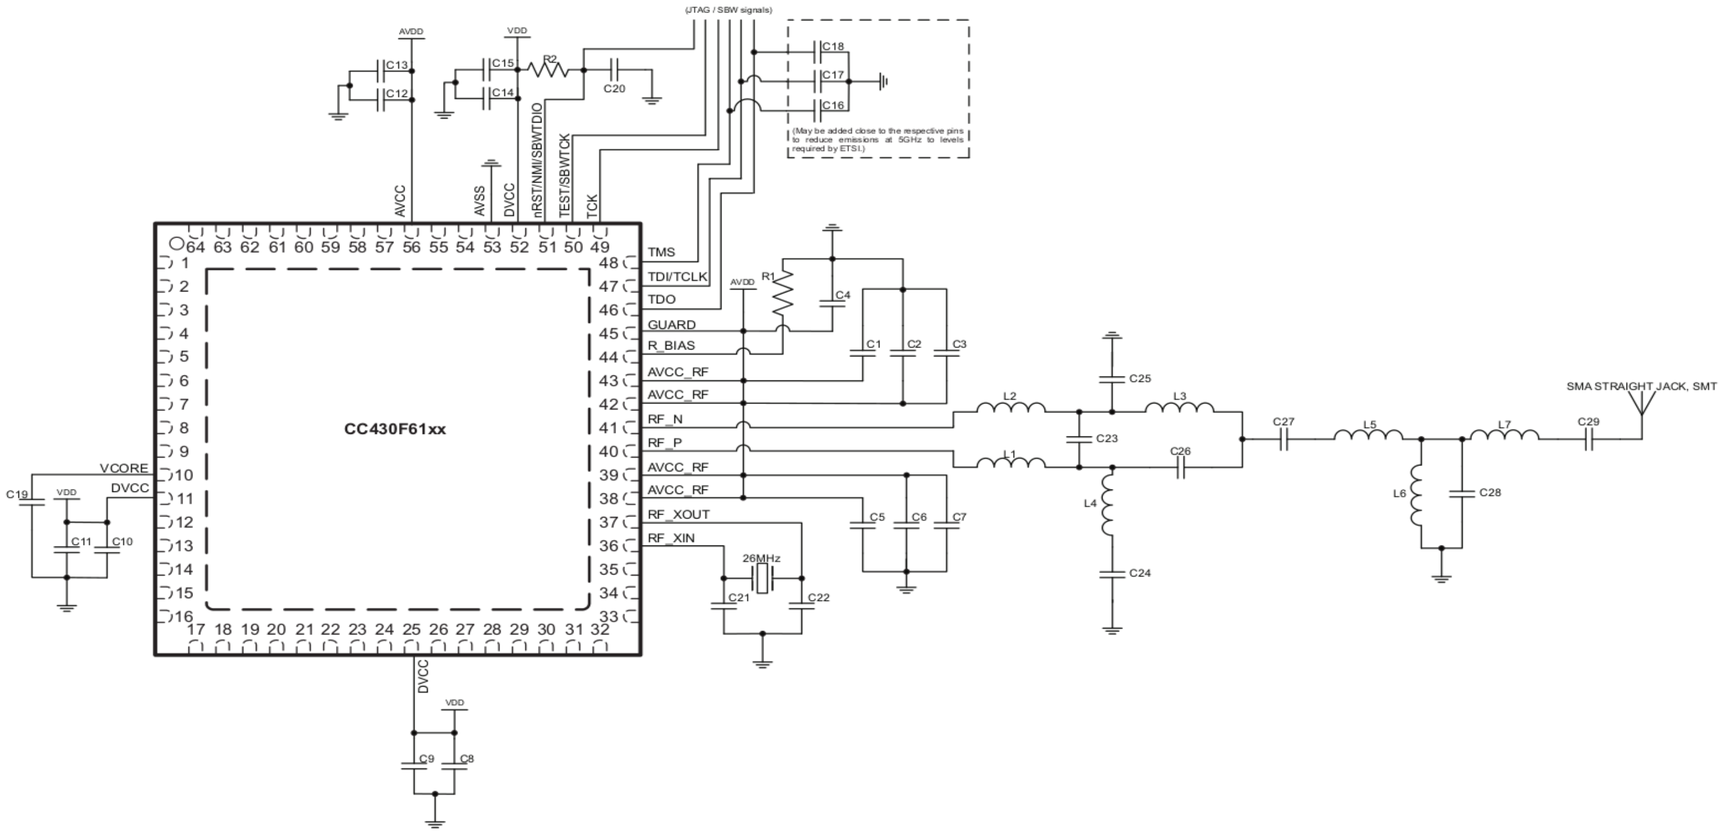

Comparison transmitting @ 1/second

- Lifetime:
  - TO-PIP: 2.2 years
  - TelosB: 3-4 months
- Size:
  - TO-PIP: 1.1 x 1.2 x 0.21 in
  - TelosB: 2.25 x 1.2 x 0.69 in
- Component cost @ quantity 1000
  - TO-PIP: $5.09
  - TelosB: $26.23
- What about channel efficiency and the receiver network?

Improving Channel Capacity

- Unslotted aloha: simply transmit packet when ready
  - Similar to TO on the transmit side
- Traditional analysis shows unslotted aloha efficiency is 18%
  - Probably of a collision grows with number of transmitters
- Do we need to sacrifice this much efficiency?
  - Carrier Sense Multiple Access (CSMA) efficiency:
  - Time-Division Multiple Access (TDMA) efficiency: 95+ %

The Capture Effect

- Radios utilize EM waves
- A stronger wave overpowers a weaker one
- Simultaneous reception of packets with different signal powers means we can recover the symbols in the stronger wave

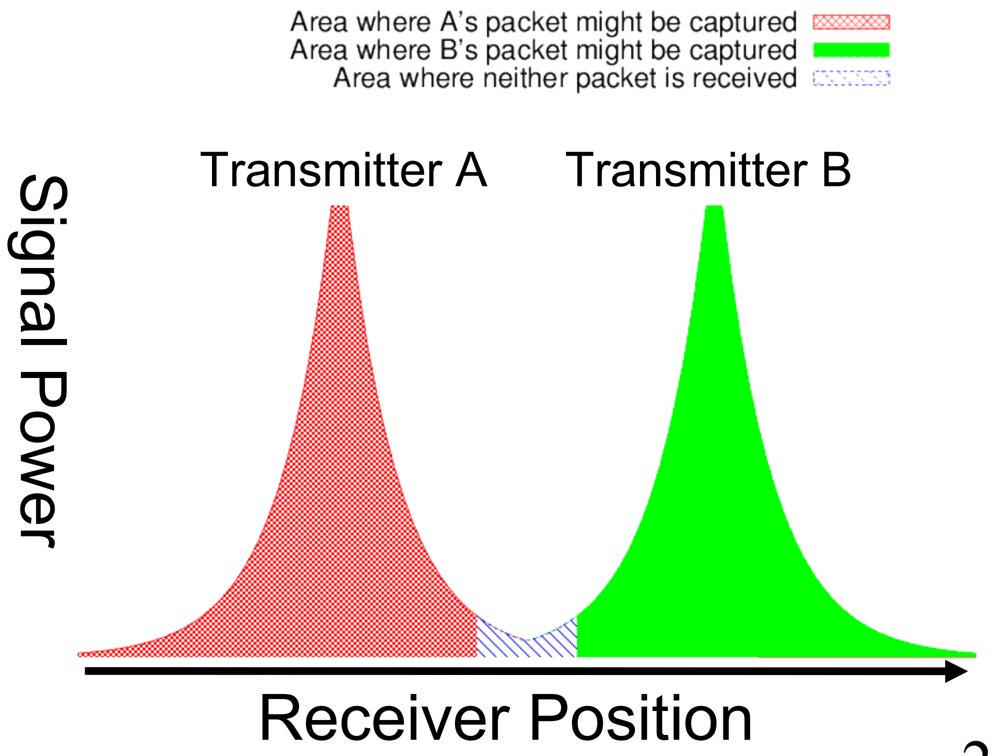

Two methods to leverage the capture effect

- Message in Messaging
  - Sense when the stronger signal arrives, and start decoding then
- Receiver Placement
  - Put receivers in physical locations where they will receive stronger signals.

Message in Messaging

Can we start receiving when the preamble of the stronger packet (pink) arrives, in spite of the interfering (grey) one?

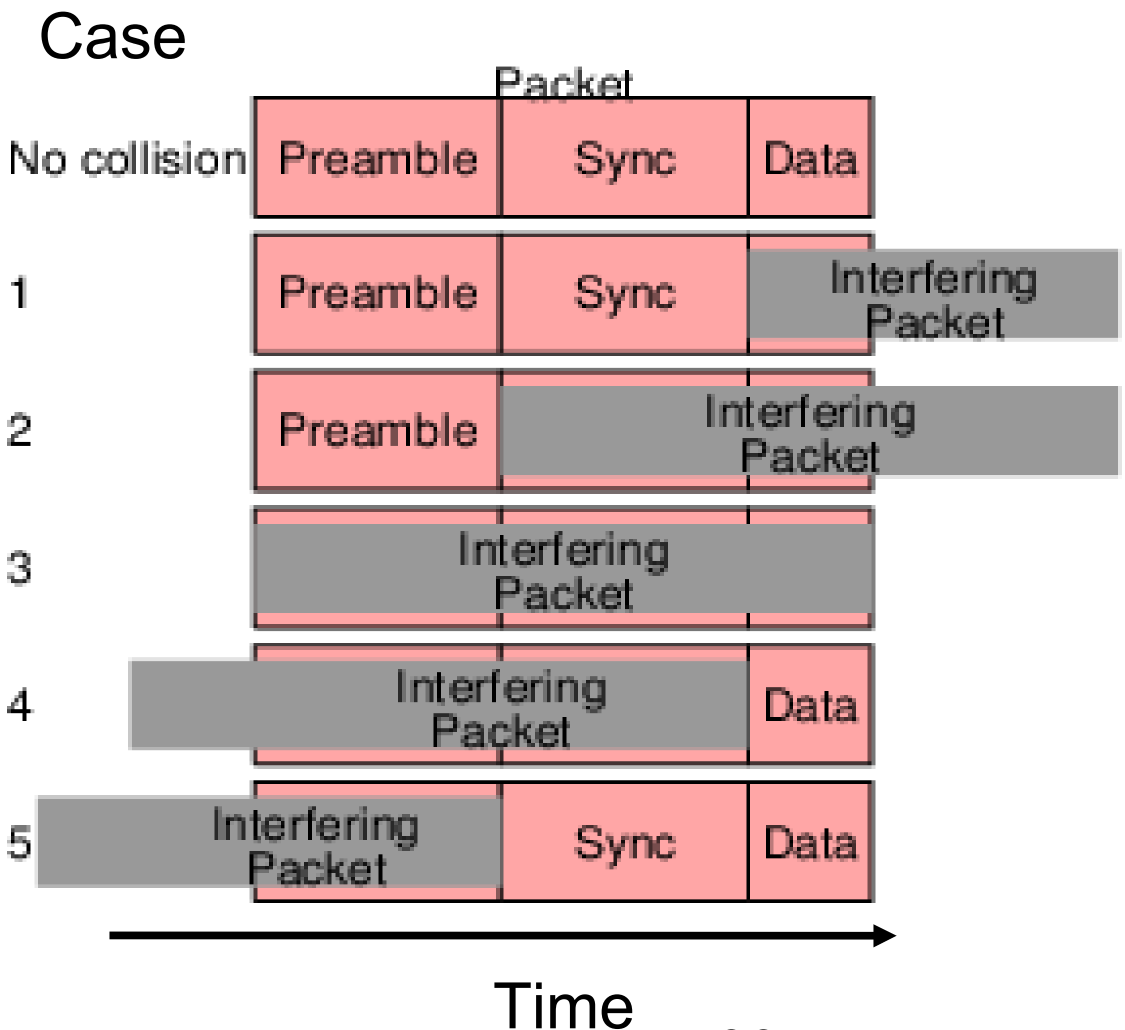

Building a MiM receiver from 2 single receivers

- If I see a preamble, tell other radio to start
  - If the second packet is stronger, the other will receive it.
- Tell other radio when I recognized a packet correctly
  - Allows aborts of a bad packet, restart to catch a new one.

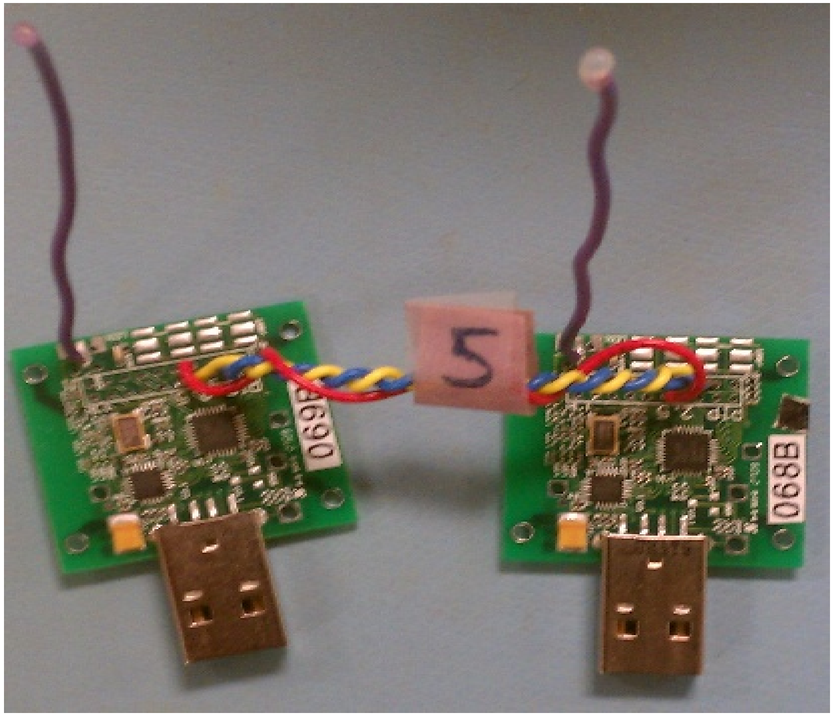

Impact of collisions – stronger packet first

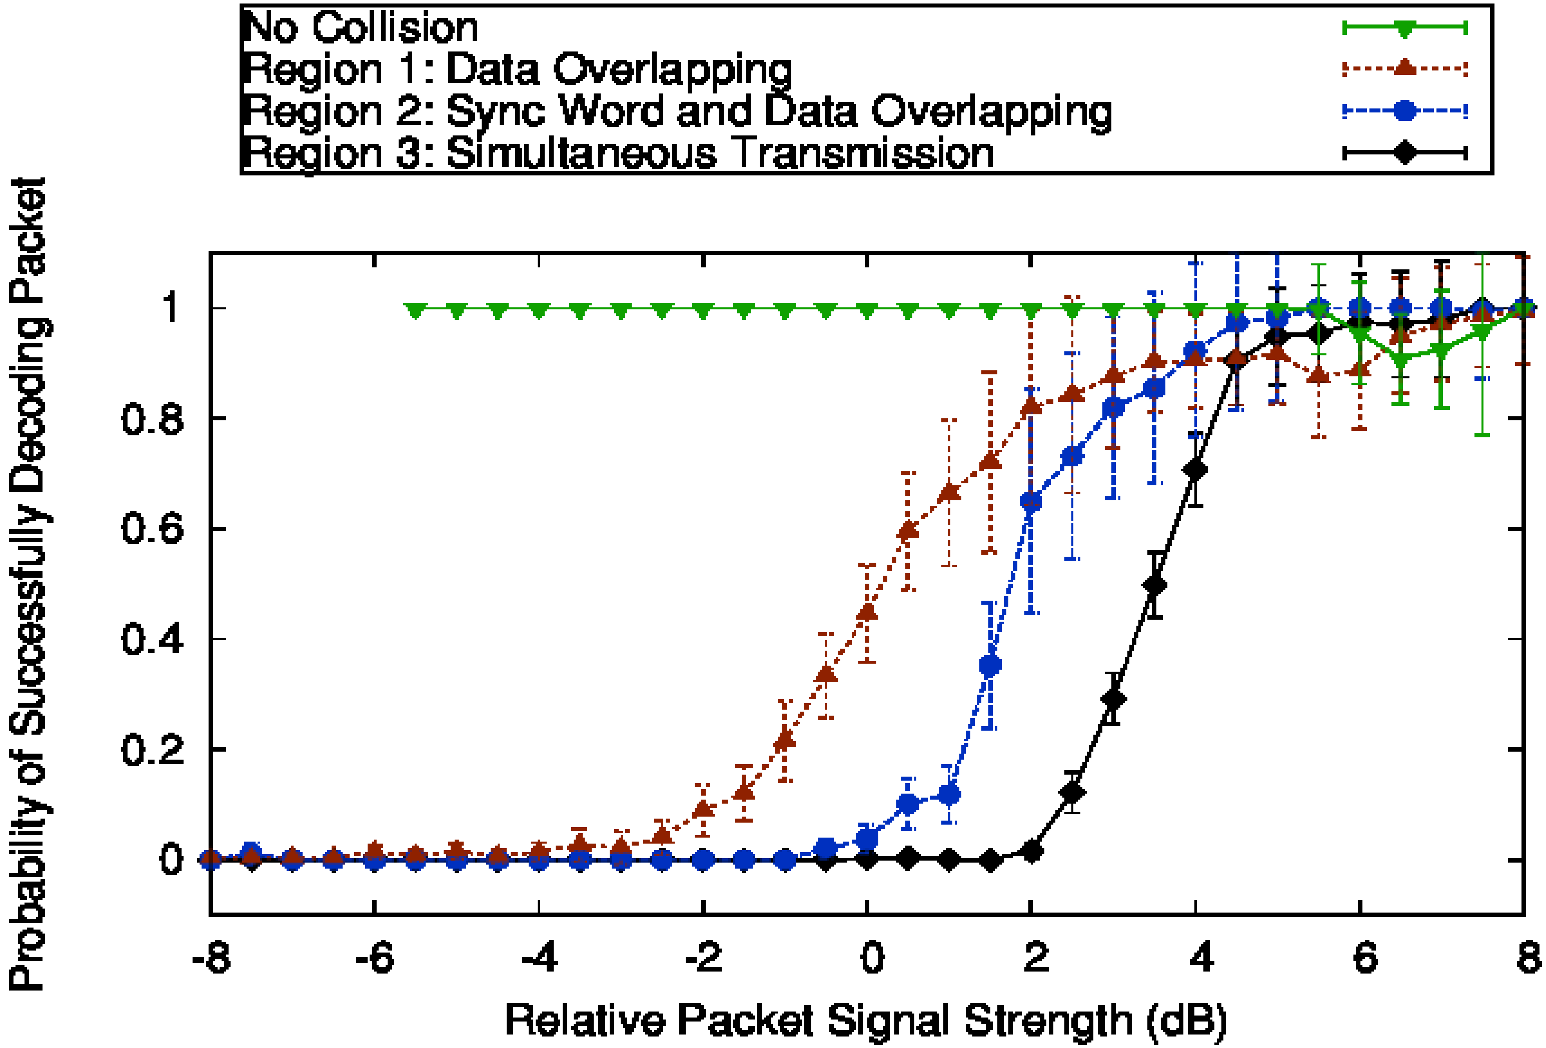

Impact of Collisions – weaker first (with MiM)

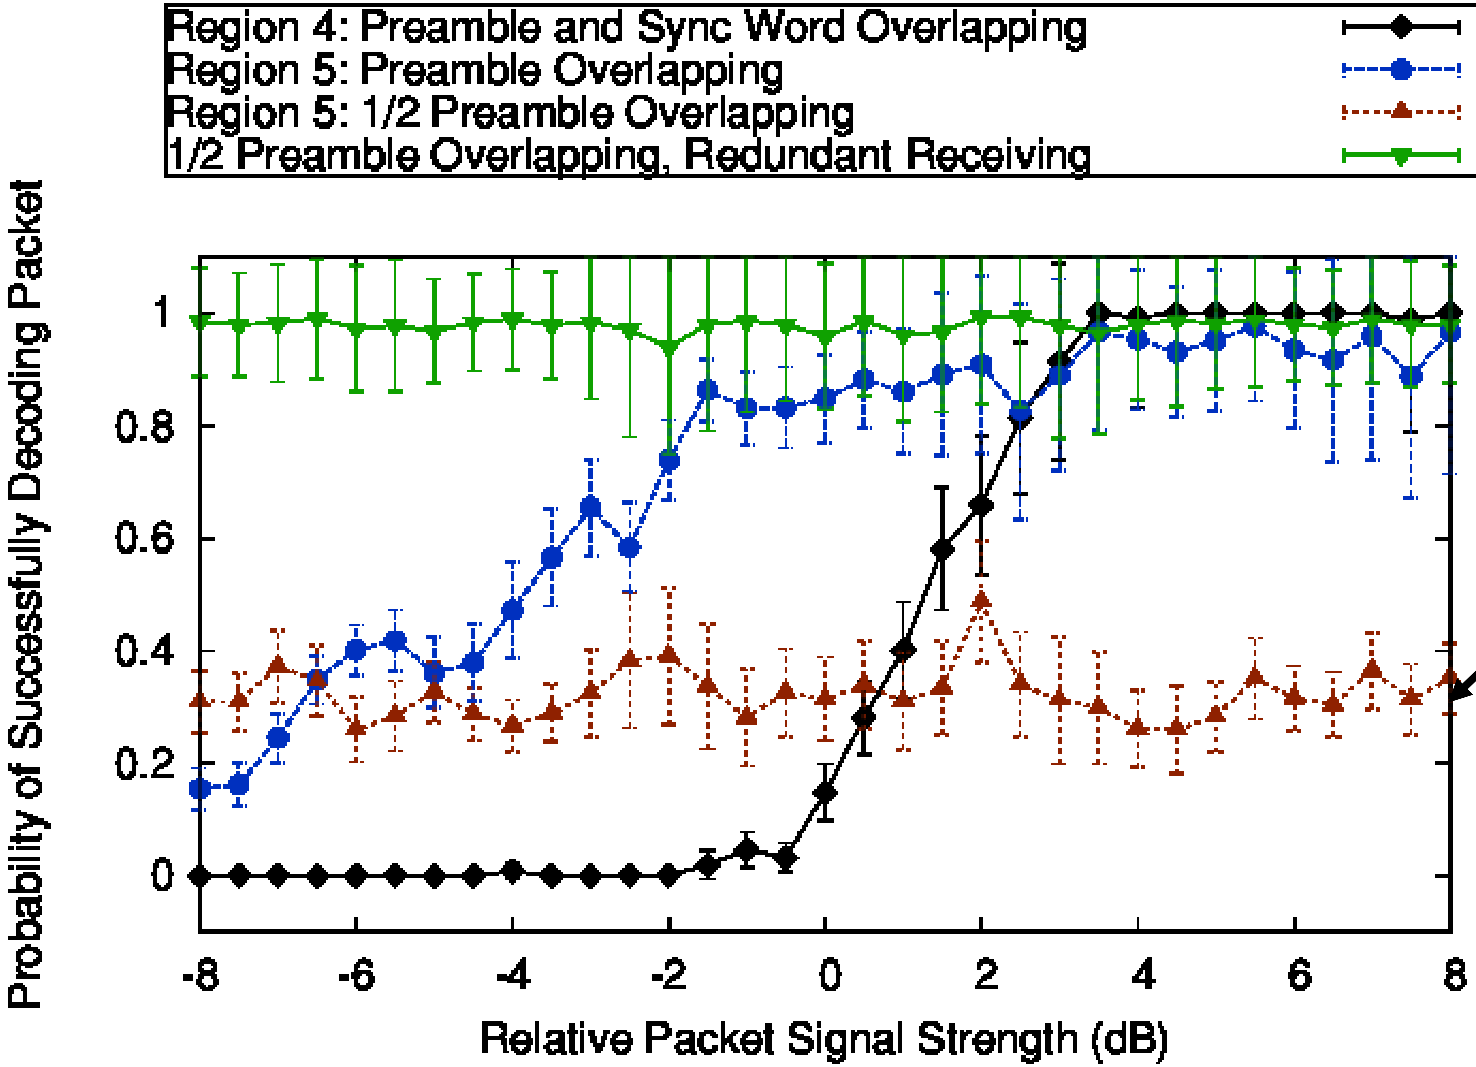

No MiM

Receiver Placement

- Given the locations of the transmitters, chose the physical locations that minimize contention
- A Capture Disk:

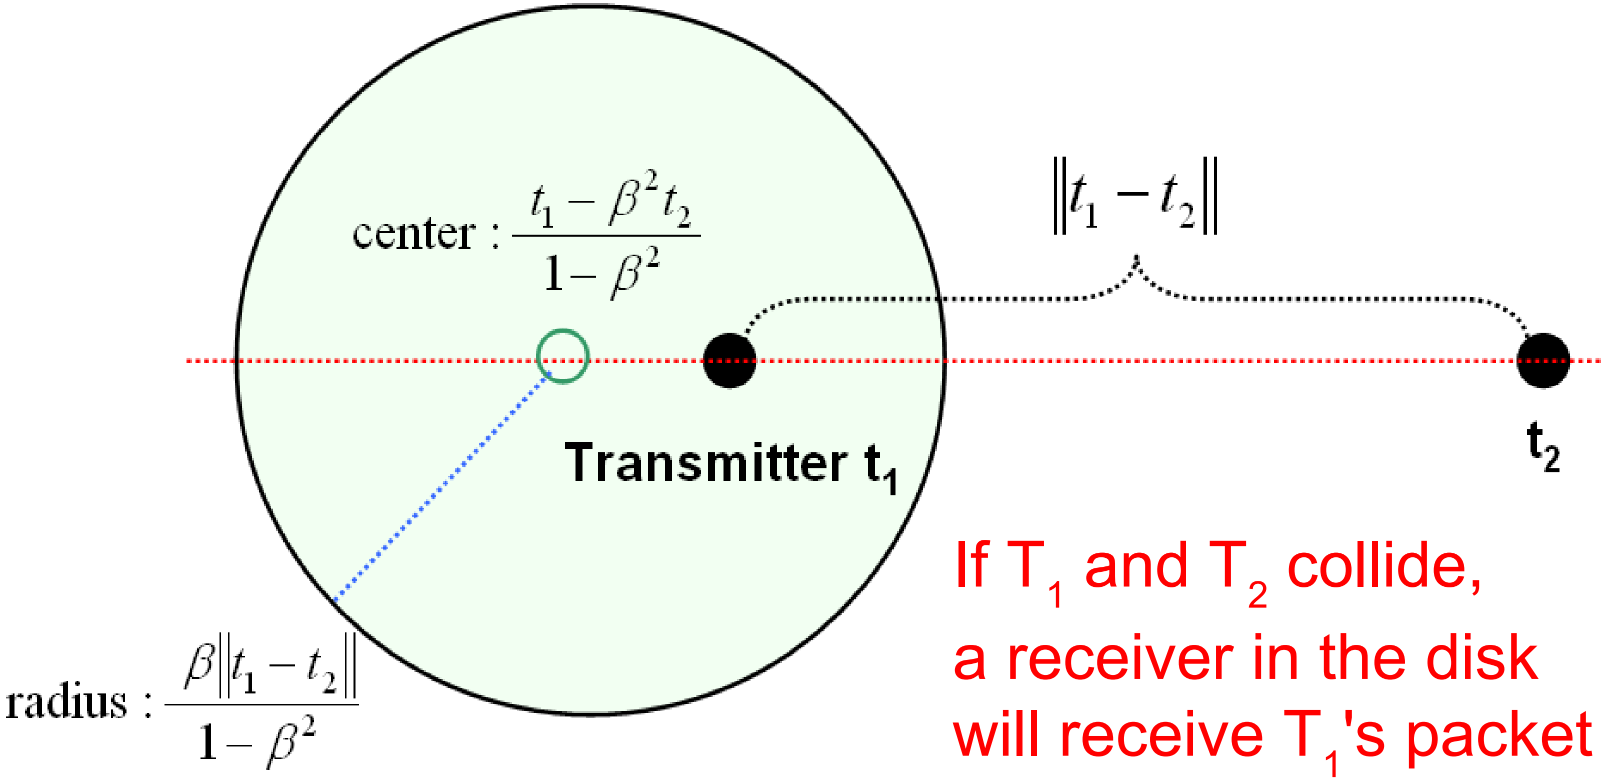

If T_1 and T_2 collide, a receiver in the disk will receive T_1's packet

F-Embed algorithm

1. Pairwise Capture Disks

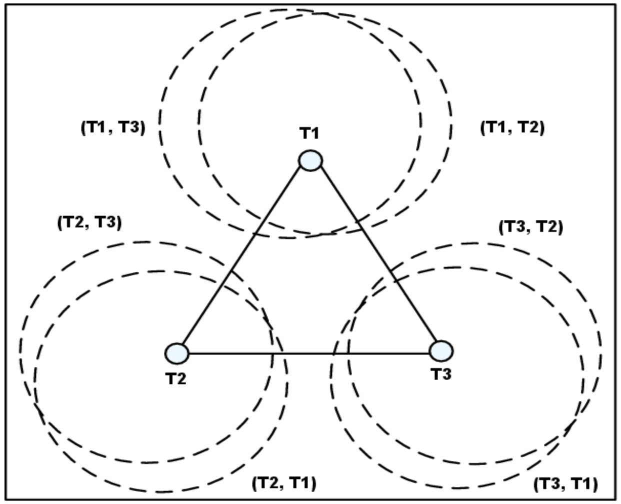

2. Possible Solution Points

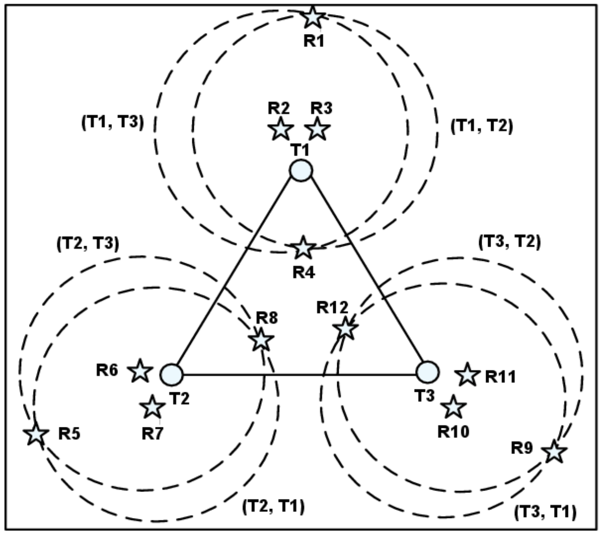

3. Bipartite Graph

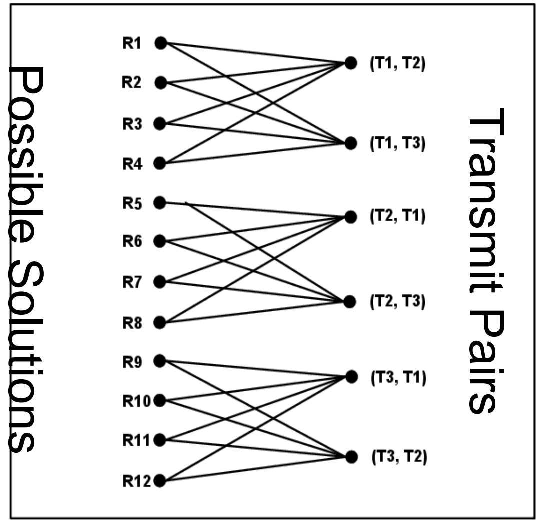

4. Top-N locations

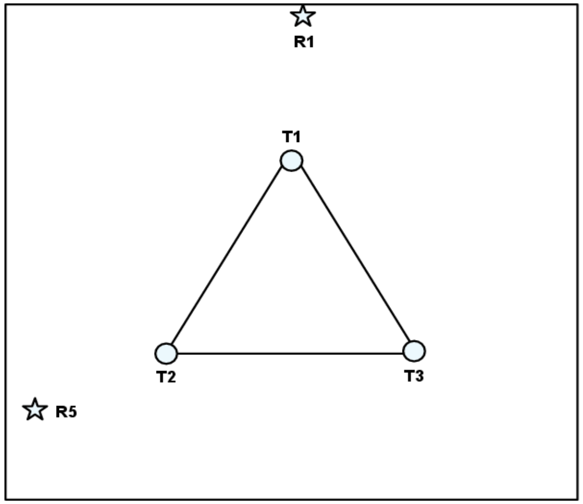

Exact solutions and approximations

- Exact solution is NP-hard
- F-Embed is 2-approximation (bounded 50% of exact)
- F-Embed is $O(R*T^6)$ in number of Receivers/Transmitters
  - Scales with number of capture disks
  - To slow for more than few 100's of transmitters
- Use a gridded approximation:
  - Divided plane into a mesh of test points (candidate solutions)
  - Scales with $C*O(n^2)$

Intuition: place receivers in the densest region

Naive Placement

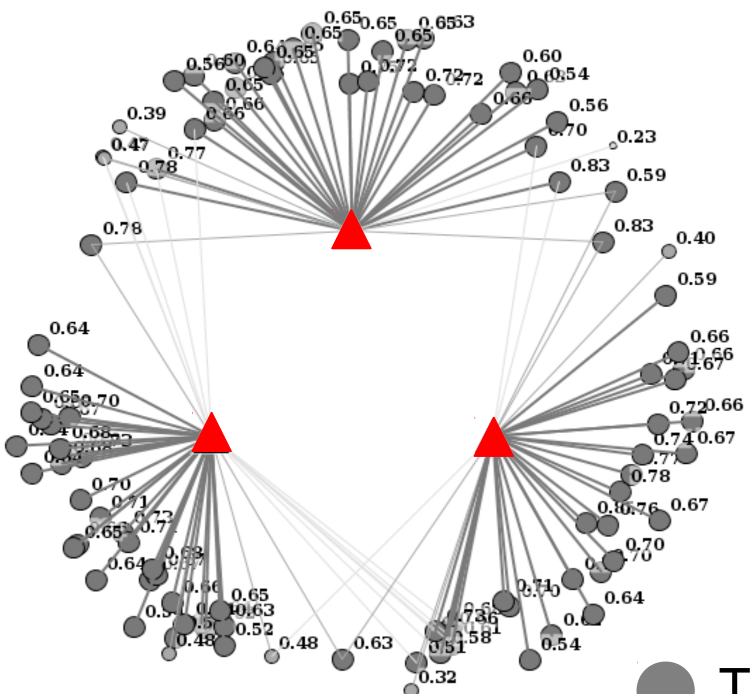

F-Embed Placement

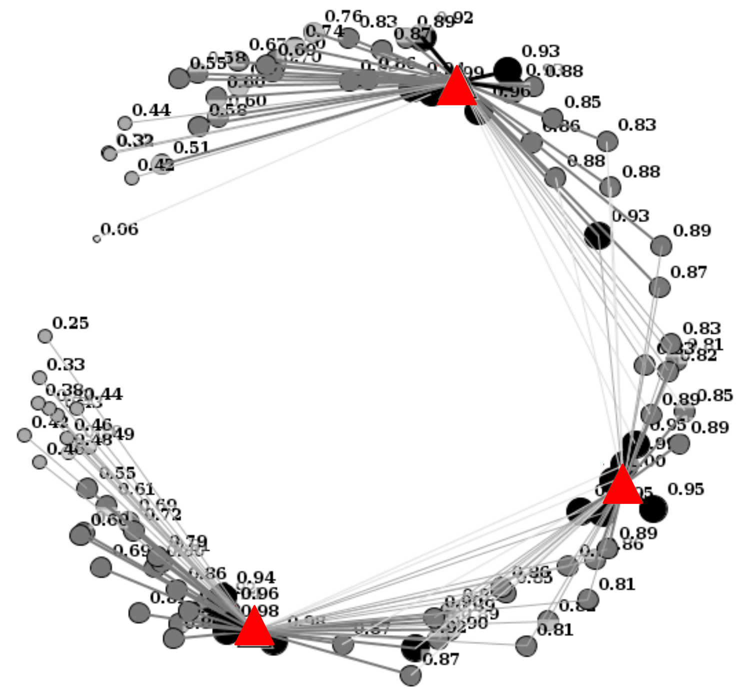

Transmitter

Receiver

Capture Probability

Channel Utilization at Scale

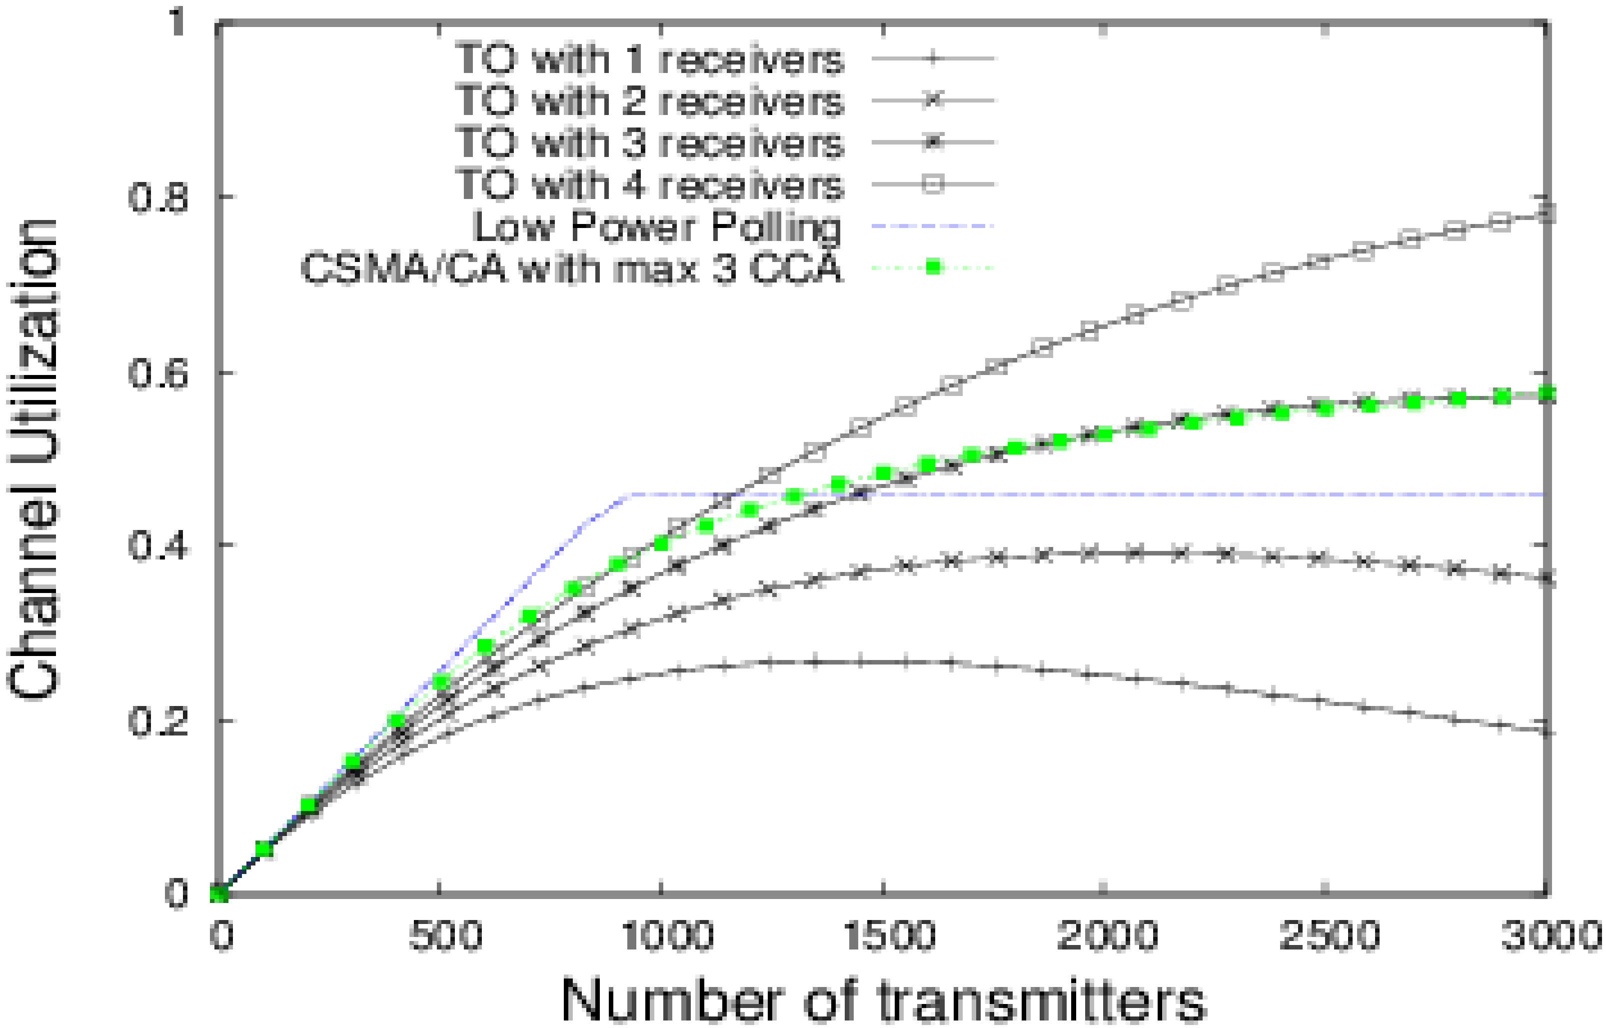

Experiment with 500 Transmitters

Bundle of 10 transmitters

MiM Receiver

Theoretic predictions vs. experiment

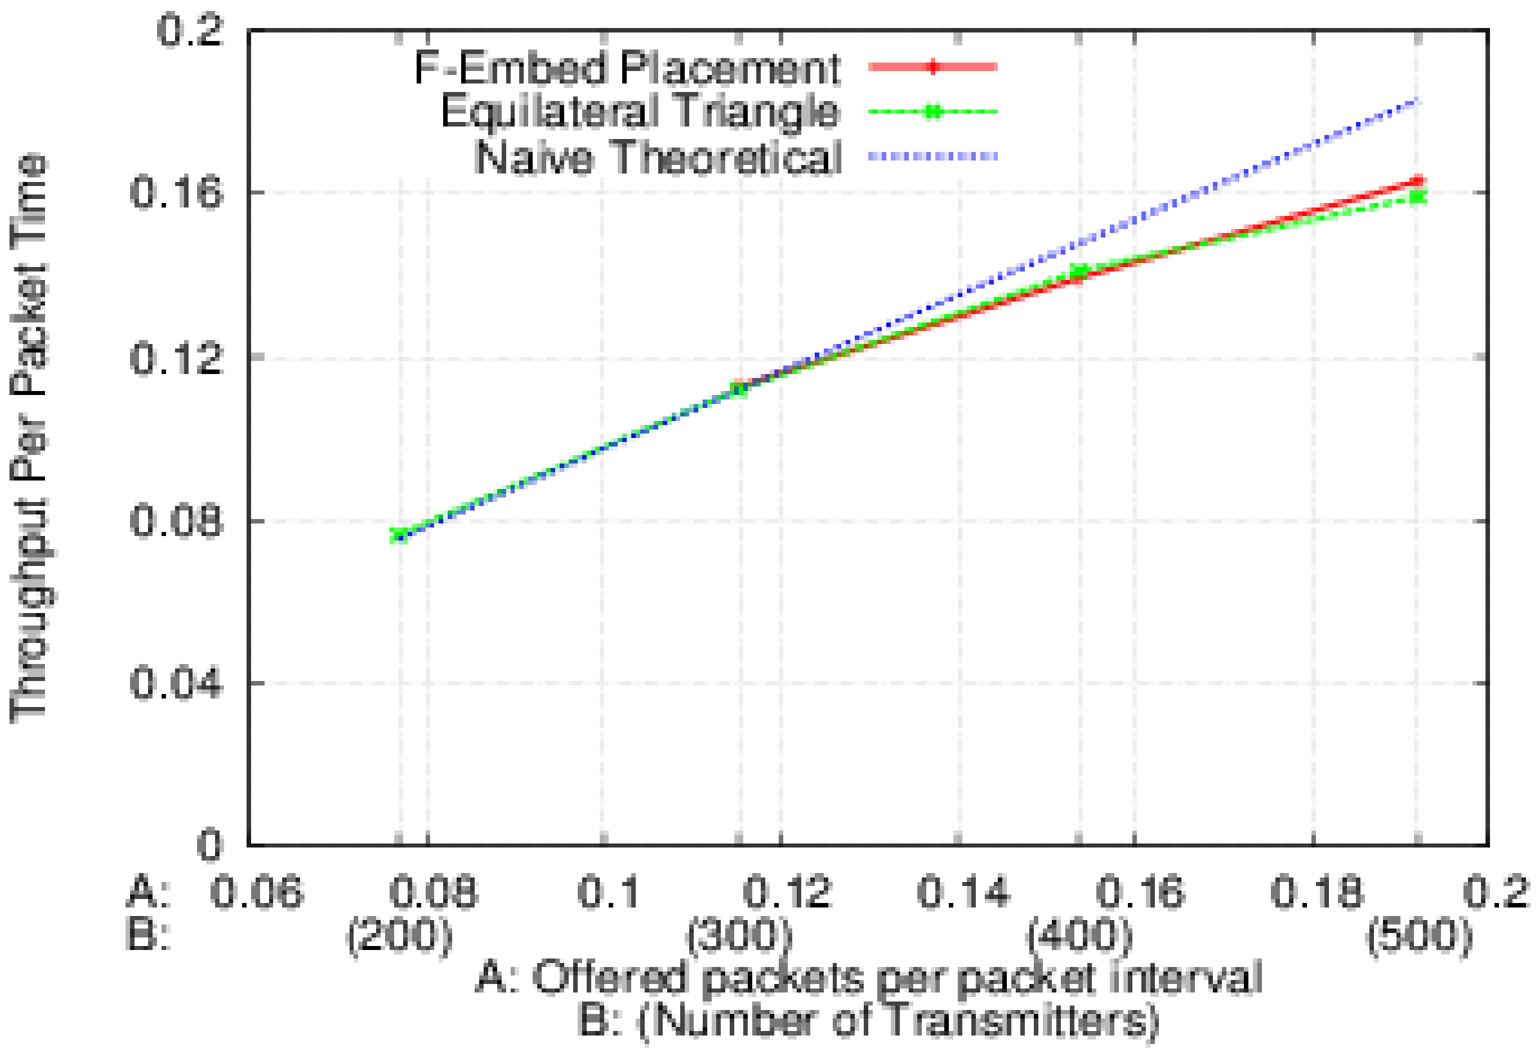

Connecting TO networks via the Owl Platform

- Sensors connect to an intermediate layer that hides details
- Solvers build higher-level representations from low-level ones
- A uniform model of the world allows sharing
- Applications run in standard environments in the cloud

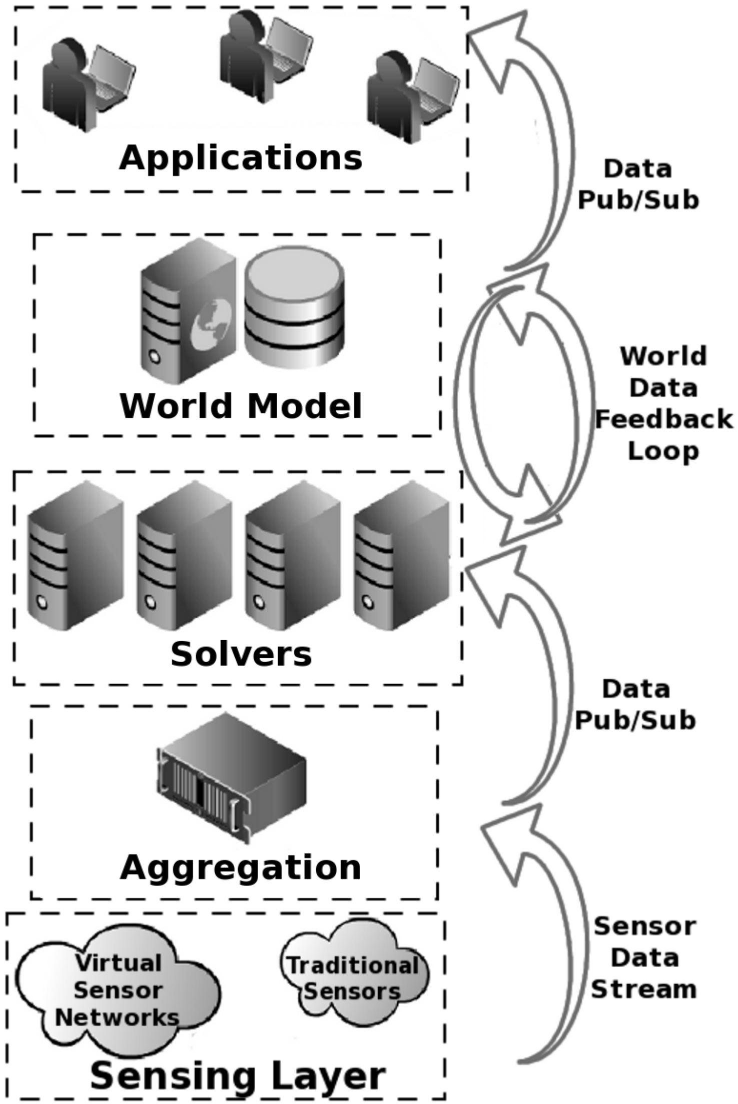

Example Applications

- Leak detection
  - Sense standing water, email/SMS if water detected
- Office space assignment
  - Sense door open/closes, assign new students to lightly used offices
- Fresh Coffee
  - Sense temperature of coffee pot, email/SMS if a temp spike
- Chair Stolen
  - Email/SMS if a chair is moved away from the owner's cubicle
- Loaner Bicycle Inventory
  - Count # of bicycles in a room to see if one is available.

Conclusions

- Channel Utilization
  - Close to 100%
- Receiver Network:
  - 1%-5% of the number of transmitters for realistic loads
  - Simple
  - Needs continuous power
  - 3-4x sensor input bandwidth
- Manageability
  - Change parameters on the sensor?
- Security
  - How to perform lightweight unidirectional security?

Conclusions

- Channel Utilization
  - Close to 100%
- Receiver Network:
  - 1%-5% of the number of transmitters for realistic loads
  - Simple
  - Needs continuous power
  - 3-4x sensor input bandwidth
- Manageability
  - Change parameters on the sensor?
- Security
  - How to perform lightweight unidirectional security?

Future Directions

- Kickstarter.com project to build base station, sensors coming soon!
- Adding a control channel:
  - Transmit mostly
- Constrained receiver placement
  - What if only specific areas (near power plugs)
- Mobile Transmitters and Receivers
- Lightweight unidirectional encryption
  - How to insure 3rd parties can't eavesdrop?
- Long data sets
  - For example, fountain codes for video streams

Backup slides

RUTGERS

Naive analytic model

Similar to simple aloha protocol models

$$P_{succ} = 1.0 - (1.0 - \frac{2\delta}{\tau})^C$$

δ=packet time (length)
τ=interval time
c=# of contenders

Modeling the Capture Effect

$$P_{succ} = 1.0 - (1.0 - \frac{2\delta}{\tau}(1 - 0.5)^M)^N$$

δ=packet time (length)
τ=interval time
0.5=chance of capture at a receiver
M=# of receivers
N=# of transmitters