

Maximizing Broadcast Throughput Under Ultra-Low-Power Constraints

Tingjun Chen¹, Javad Ghaderi¹, Dan Rubenstein², and Gil Zussman¹

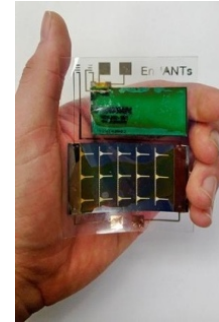
¹Electrical Engineering and ²Computer Science, Columbia University

Dec. 15, 2016



Tag-to-Tag Networks

- Internet-of-Things (IoT)
- Small, flexible, and energetically self-reliant tags



Smart Supermarket



Supply Chain Management

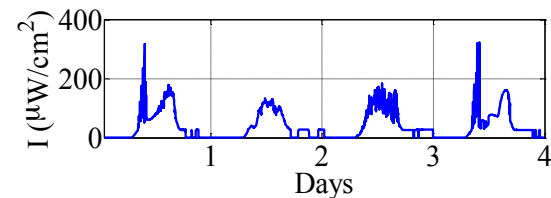
- Enabling technologies:
 - Energy harvesting
 - Ultra-low-power wireless communication

New Challenges

- Tags utilize ultra-low-power wireless communication
 - Power consumption: nJ/bit

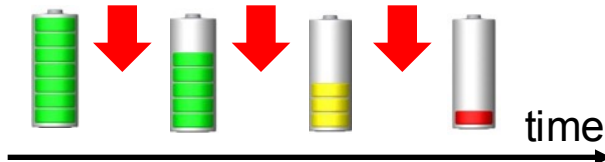


- Tags can harvest and store energy
 - Power budget is $\sim 10 \mu\text{W}$

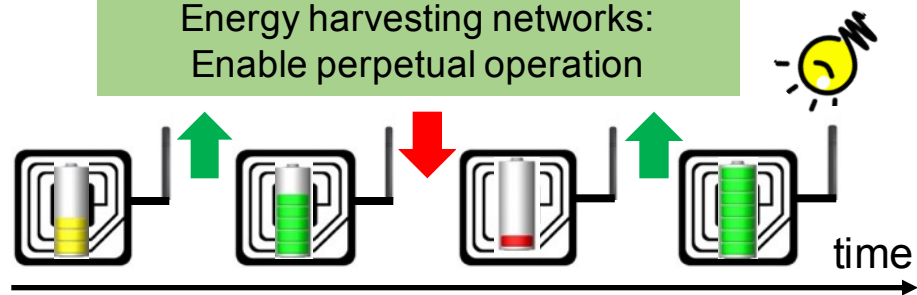


Indoor light energy measurement

Classical sensor networks:
Maximize lifetime



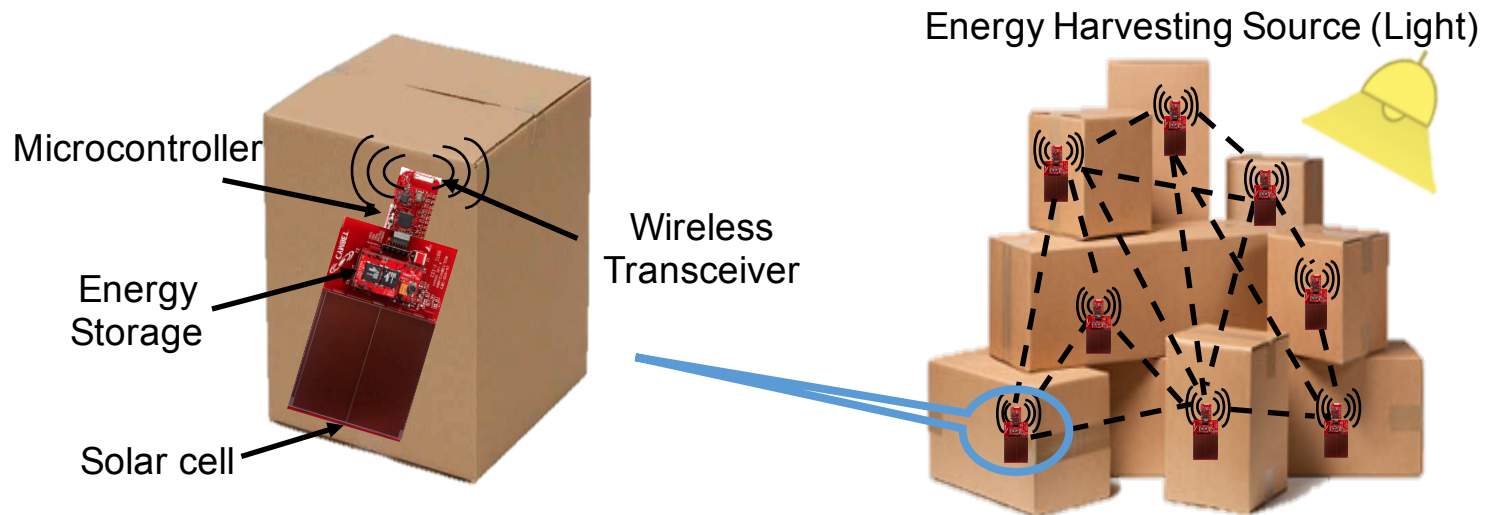
Energy harvesting networks:
Enable perpetual operation



- However, energy availability limits the tags' capabilities
 - Low data rate (e.g., 3 Kbps with indoor lights), limited control information
 - No accurate clocks for synchronization
 - Heterogeneous networks: tags have different power budgets and TX/RX power

The Application and Objective

- An example application: Locating misplaced boxes in a warehouse
 - Boxes equipped with active tags perform neighbor discovery
 - Related work: [Gorlatova et al. *MobiCom*'09], [Liu et al. *SIGCOMM*'13], [Wang et al. *SIGCOMM*'13], [Margolies et al. *ToSN*'15], etc.



1. Obtain the **maximum achievable throughput** in a network of heterogeneous ultra-low-power nodes with given power budgets.
2. Develop an asynchronous protocol that achieves the maximum throughput in a **distributed manner**.

Outline

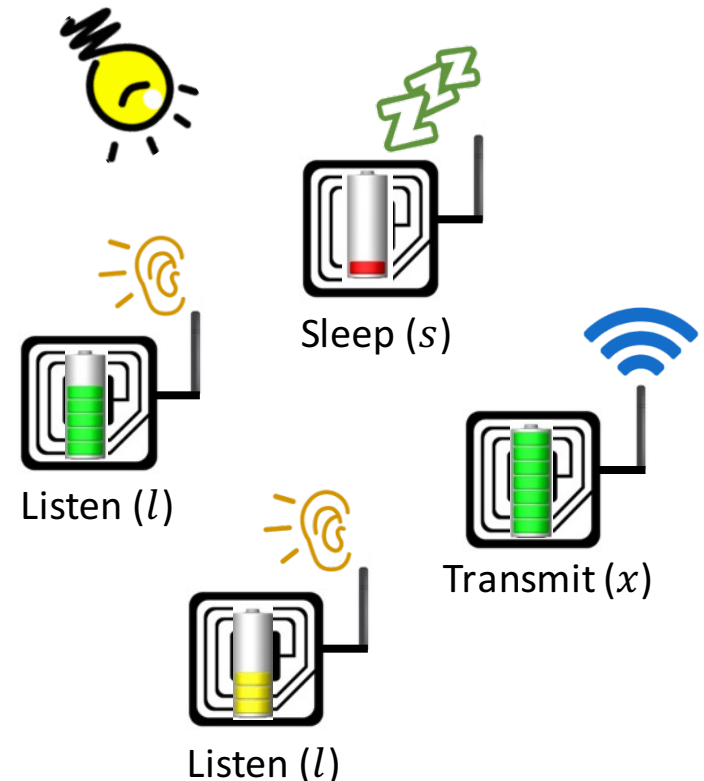
- Energy-constrained Broadcast Throughput Maximization Problem
 - **Goal**: Maximize broadcast throughput between ultra-low-power nodes
 - **Subject to**: all heterogeneous nodes maintain within their power budgets
- *EconCast* (**E**nergy-**con**strained Broad**Cast**) – A simple asynchronous distributed protocol that provably approaches the maximum throughput
- Performance evaluation of *EconCast* via simulations and experimentation

Related Work

- Low-power Medium Access Control (MAC) in sensor networking
 - S-MAC [Ye et al. 2002], T-MAC [Van Dam et al. 2003], WiseMAC [A El-Hoiydi et al. 2004], B-MAC [Polastre et al. 2006], RI-MAC [Sun et al. 2008], A-MAC [Dutta et al. 2010], etc.
 - *Synchronous MAC*: broadcast schedule. *Asynchronous MAC*: send long preambles
- Neighbor discovery in low-power wireless networks
 - Birthday [McGlynn et al. 2001], Disco [Dutta et al. 2008], U-Connect [Kandhalu et al. 2010], Searchlight [Bakht et al. 2012], Hello [Sun et al. 2014], Panda [Margolies et al. 2016], etc.
 - *Deterministic*: guaranteed worst case latency. *Probabilistic*: higher throughput
 - Global knowledge is known a priori; nodes are homogeneous in a network
- Q-CSMA and utility maximization
 - Utility-optimal CSMA [Lee et al. 2009], Distributed CSMA [Jiang et al. 2010], Multi-hop CSMA [Xu et al. 2010], Q-CSMA [Ni et al. 2012], etc.
 - Queues buffer “energy”, not “data”

Problem Formulation

- N nodes in a heterogeneous clique network
- Each node i has a power budget ρ_i (uW), and can be in 3 states:
 - sleep (s) consumes 0 power
 - listen (l) consumes L_i (uW) power
 - transmit (x) consumes X_i (uW) power
- Severe energy constraints: power budget is much lower than listen/transmit power
- Network state: $\mathbf{w} \in \{s, l, x\}^N$
- Network state throughput: $T_{\mathbf{w}}$



Listen power: L_i
Transmit power: X_i
Network state: $\mathbf{w} = [l, l, s, x]$
Throughput of state \mathbf{w} : $T_{\mathbf{w}} = 2$

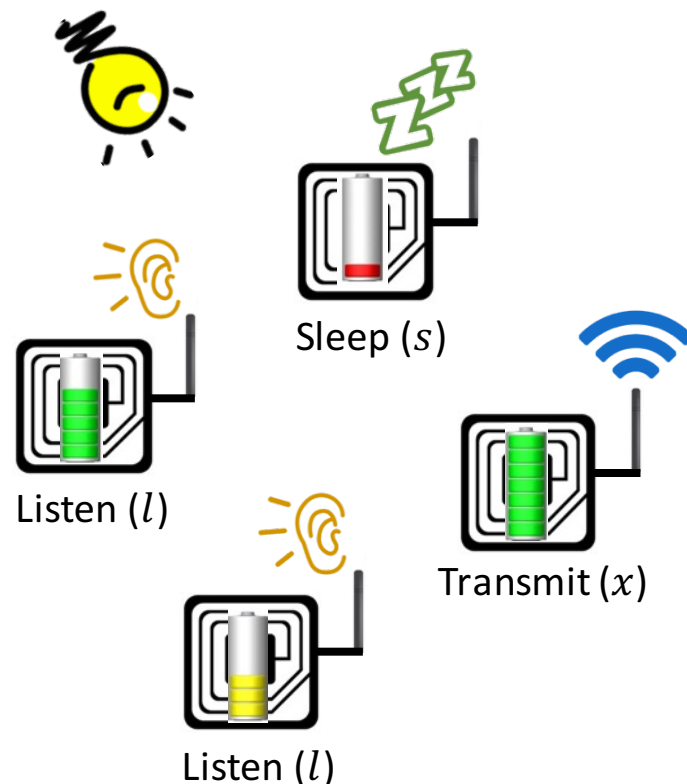
Problem Formulation

- $\pi_{\mathbf{w}}$: fraction of time the network is in state \mathbf{w}
- α_i : fraction of time node i listens
- β_i : fraction of time node i transmits
- $\alpha_i = \sum_{\mathbf{w}: i \text{ listen}} \pi_{\mathbf{w}}$
- $\beta_i = \sum_{\mathbf{w}: i \text{ transmit}} \pi_{\mathbf{w}}$
- Energy-constrained Broadcast Throughput Maximization Problem

$$\max_{\mathbf{w}} : \sum_{\mathbf{w}} T_{\mathbf{w}} \pi_{\mathbf{w}}$$

$$\text{s.t.: } \alpha_i L_i + \beta_i X_i \leq \rho_i, \quad \forall i \quad \leftarrow \text{Energy constraint}$$

$$\sum_{\mathbf{w}} \pi_{\mathbf{w}} = 1, \quad \pi_{\mathbf{w}} \geq 0, \quad \forall \mathbf{w}$$



Listen power: L_i

Transmit power: X_i

Network state: $\mathbf{w} = [l, l, s, x]$

Throughput of state \mathbf{w} : $T_{\mathbf{w}} = 2$

A Simple Example

- 4 nodes in a clique with equal listen/transmit power: $L_i = X_i = 100 \text{ uW}, \forall i$
- Optimal **centralized** solution:

	Homogeneous				Heterogeneous			
	1	2	3	4	1	2	3	4
Power budget ρ_i (uW)	10	10	10	10	50	10	5	1
Awake time (%)	10	10	10	10	50	10	5.0	1.0
OPT listen fraction α_i (%)	7.5	7.5	7.5	7.5	7.4	4.6	3.1	0.9
OPT transmit fraction β_i (%)	2.5	2.5	2.5	2.5	42.6	5.4	1.9	0.1

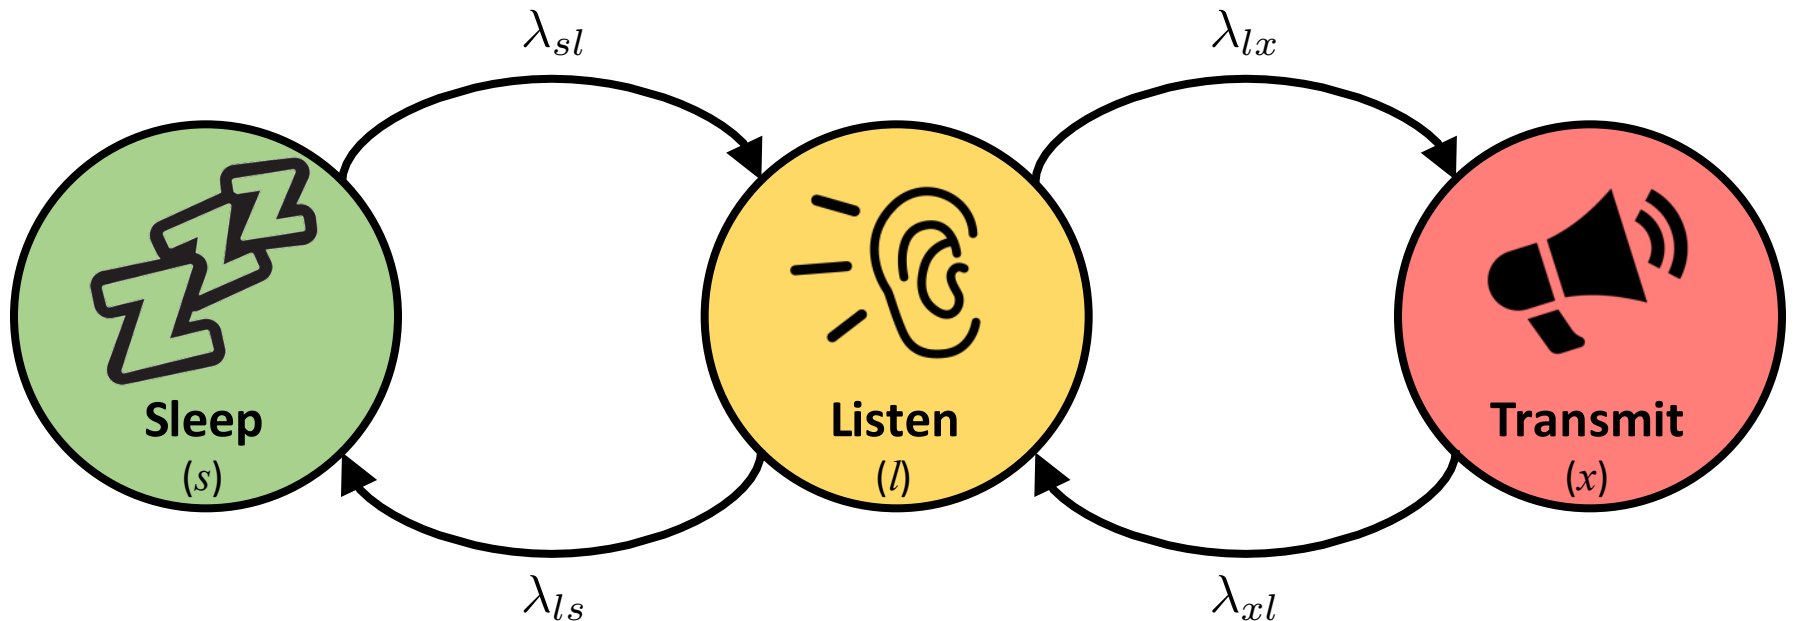
The change of other nodes' properties **changes** the “behavior” of node 2

Goal: Achieve the maximum throughput in a **distributed manner**

Solution: *EconCast* – an asynchronous distributed protocol

EconCast: Design

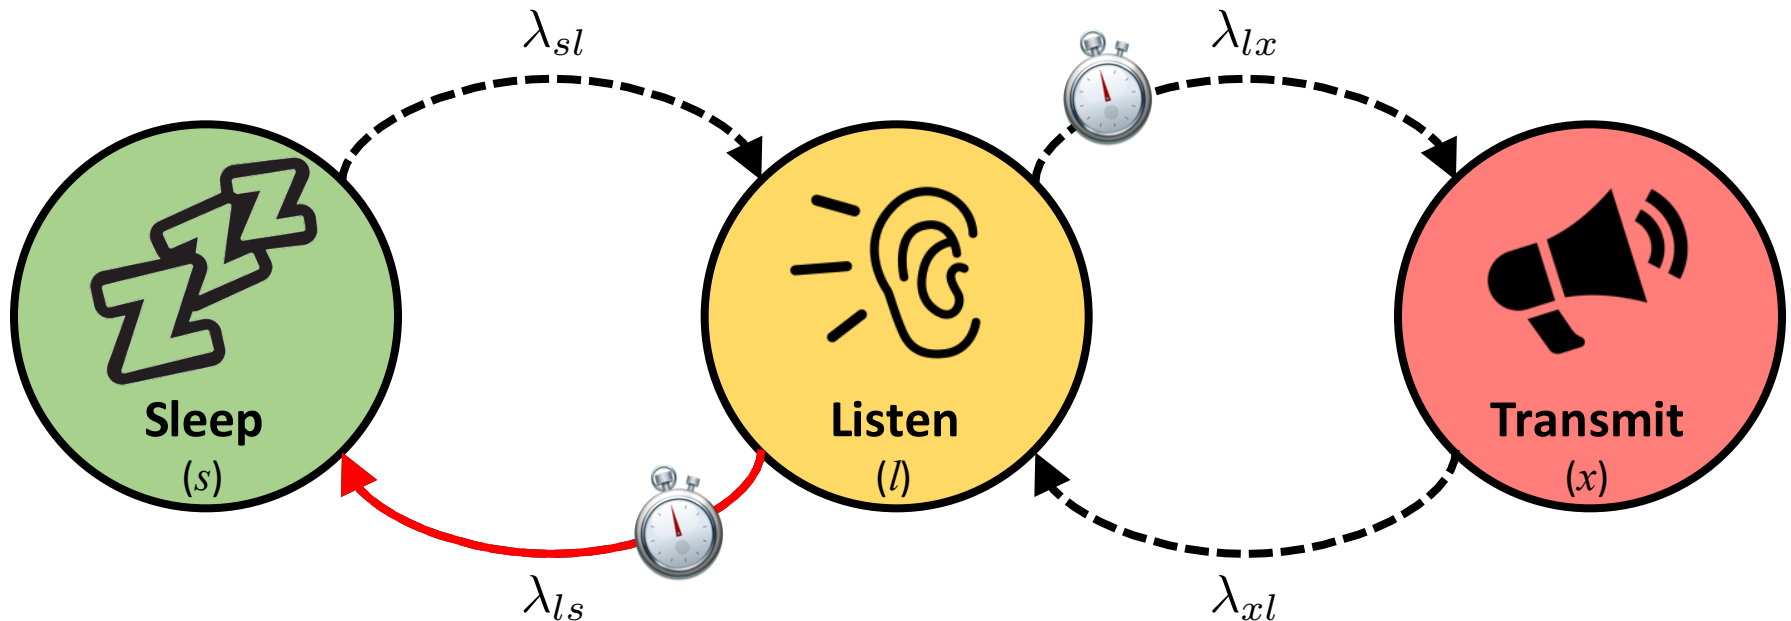
- Transition between sleep, listen, and transmit states
- The duration in each state is exponentially distributed
- Perform carrier-sensing when waking up



Transition for a single node

EconCast: Design

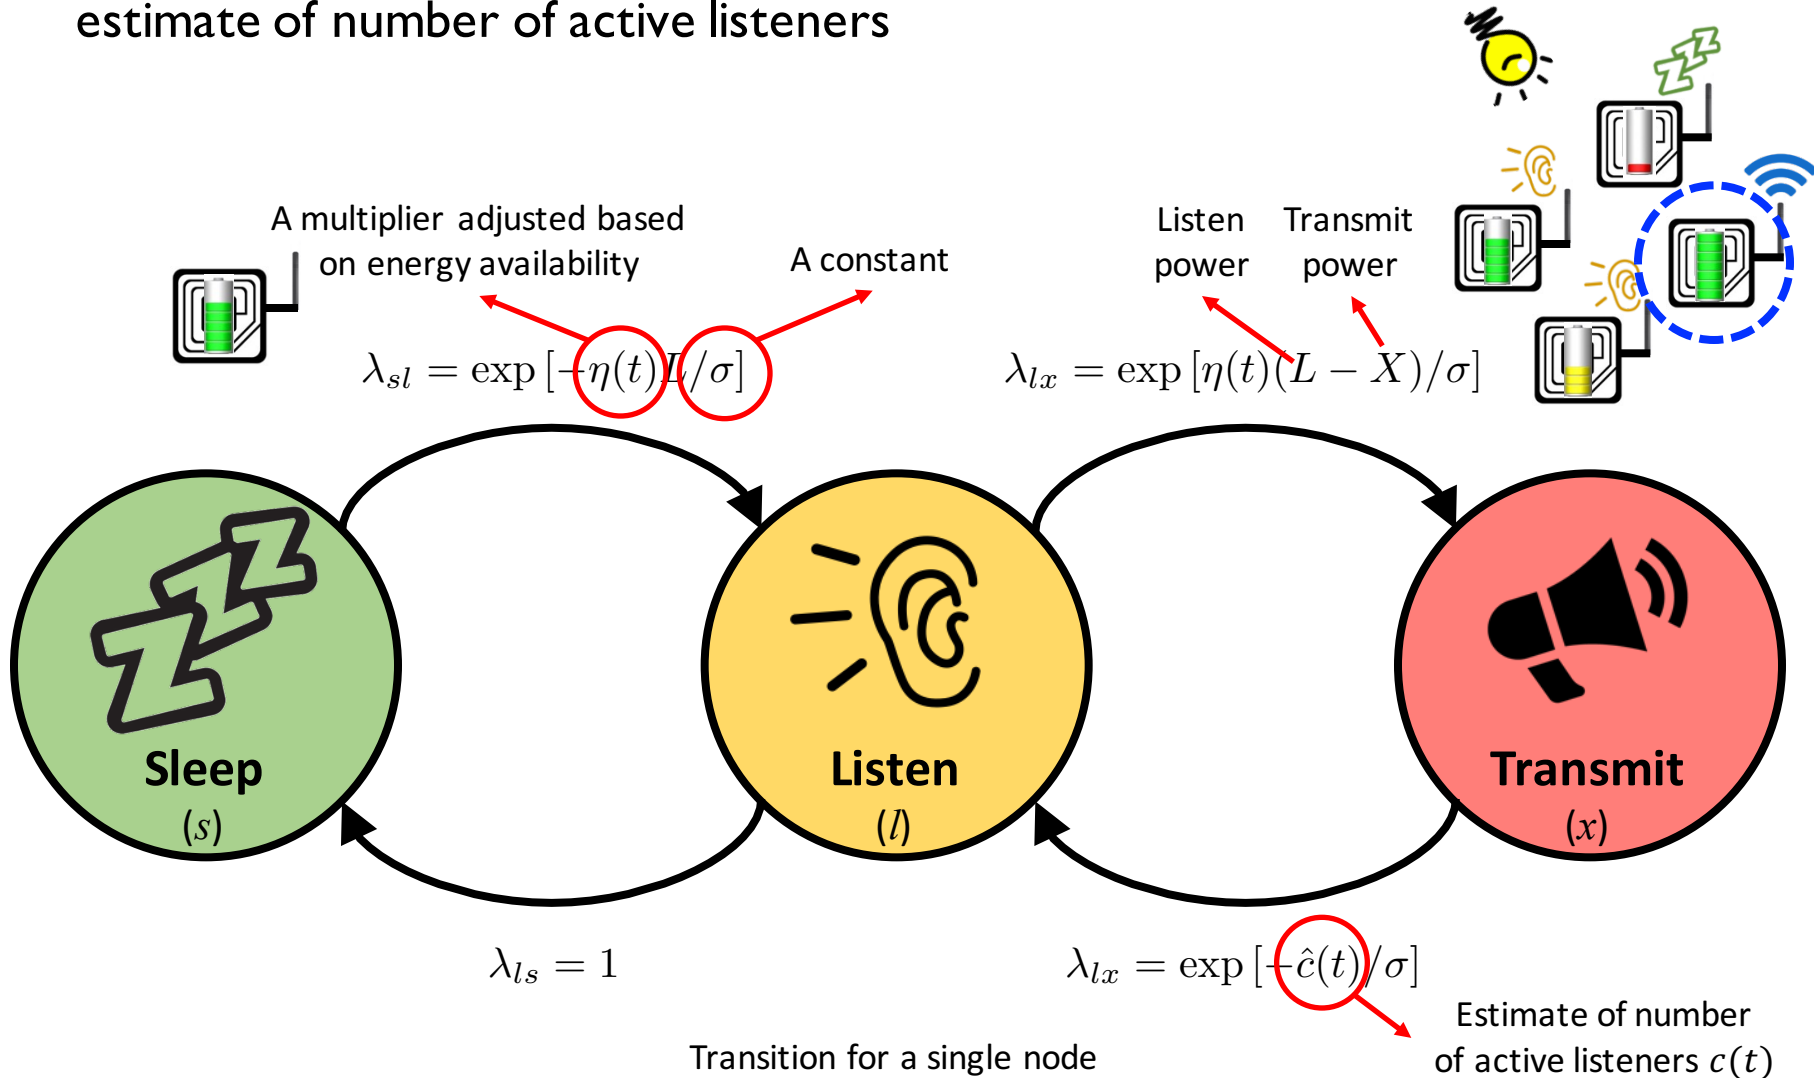
- Transition between sleep, listen, and transmit states
- The duration in each state is exponentially distributed
- Perform carrier-sensing when waking up



A node enters listen state...

EconCast: Design

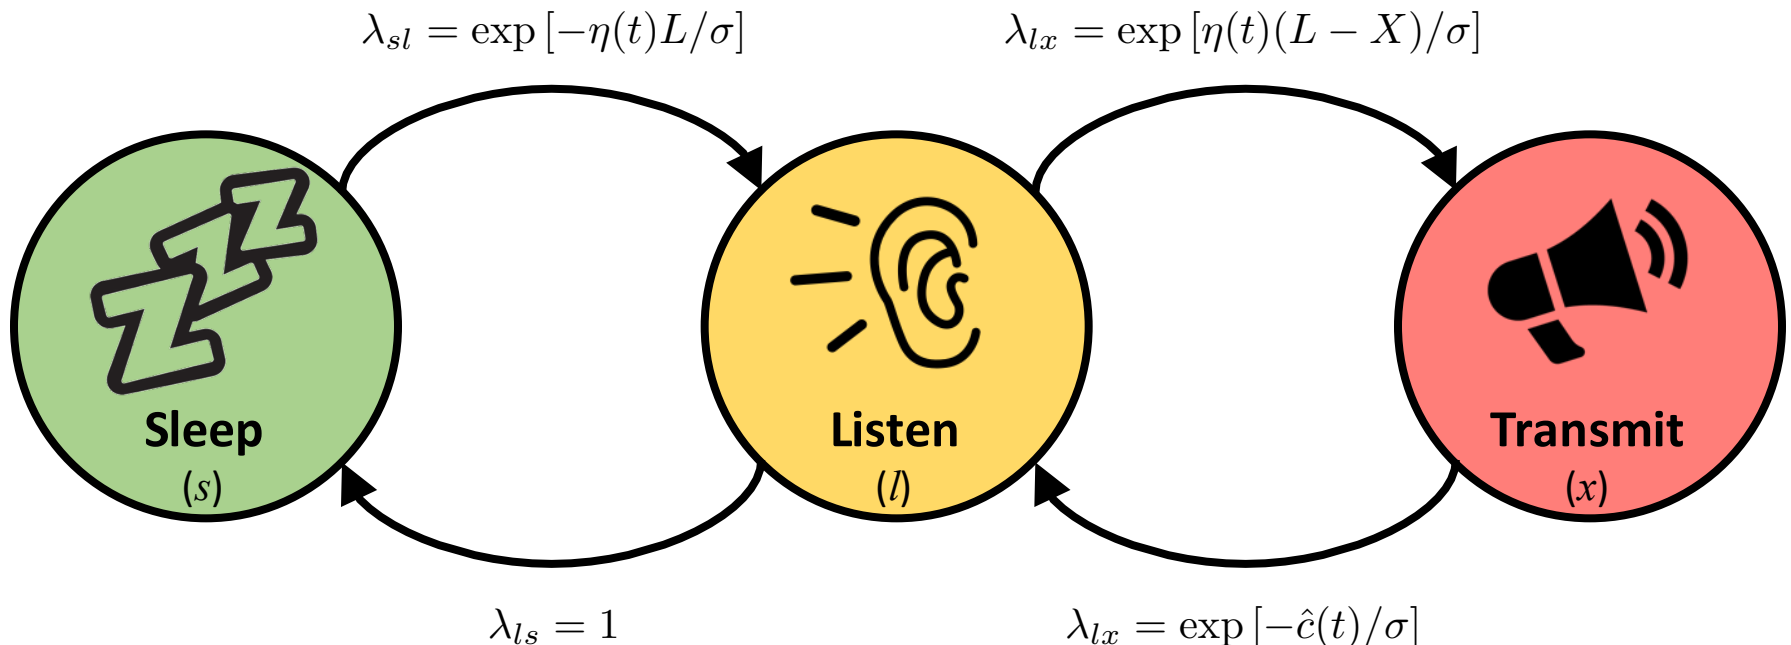
- Transition rates are adjusted over time t based on energy availability and estimate of number of active listeners



EconCast: Analysis

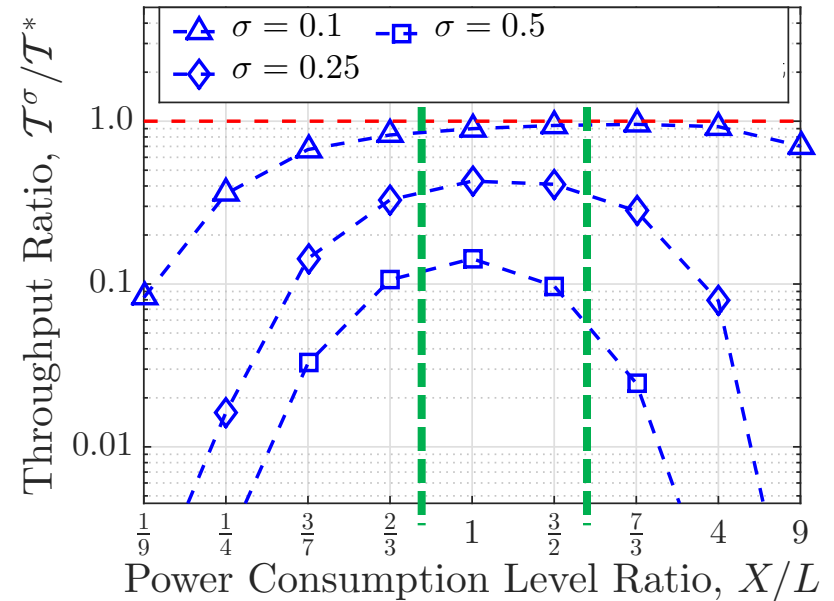
Theorem 1: Let $\sigma \rightarrow 0$ and under perfect knowledge of the number of active listeners (i.e., $\hat{c}(t) = c(t)$), the average throughput of **EconCast** converges to the maximal achievable throughput.

Proof: Based on the Markov Chain Monte Carlo approach.



Simulation Evaluation – Throughput

- 5 nodes, clique topology, homogeneous nodes, vary X/L
- $\rho = 10 \text{ uW}$, $L + X = 1000 \text{ uW}$
- Ratio between throughput achieved by **EconCast** and the maximum throughput
- Also compare to
 - Panda [INFOCOM'16]
 - Searchlight [MobiCom'12]
 - Birthday [MobiHoc'01]



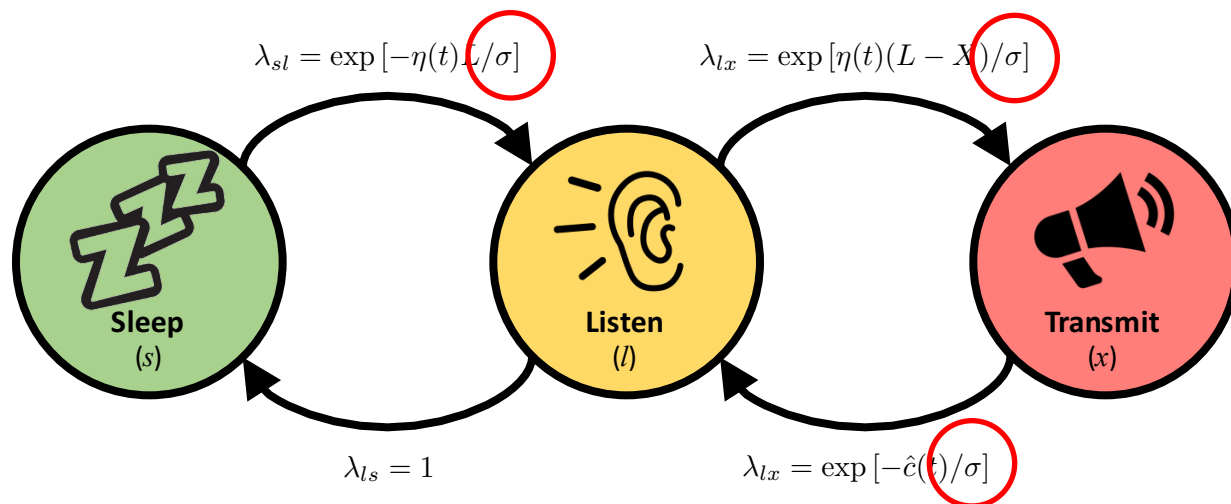
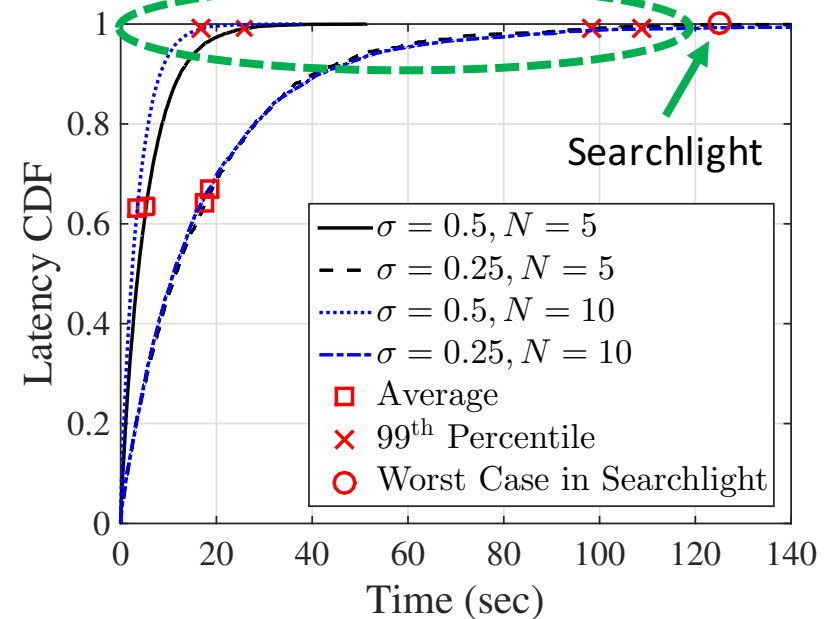
TI Radio	Listen power L (mW)	Transmit power X (mW)	Ratio X/L
CC2541	59 – 67	55 – 60	0.8 – 1.0
CC2500	65	59 – 75	0.9 – 1.2
CC2640	19	21 – 30	1.1 – 1.6

- Throughput drops with increased σ
- At $L = X$, **EconCast** outperforms Panda by 6x ($\sigma = 0.5$) and 17x ($\sigma = 0.25$)

Simulation Evaluation – Effect of σ

- CDF of latency (i.e., time between successful transmissions)

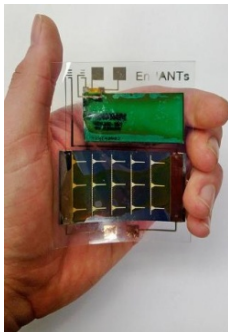
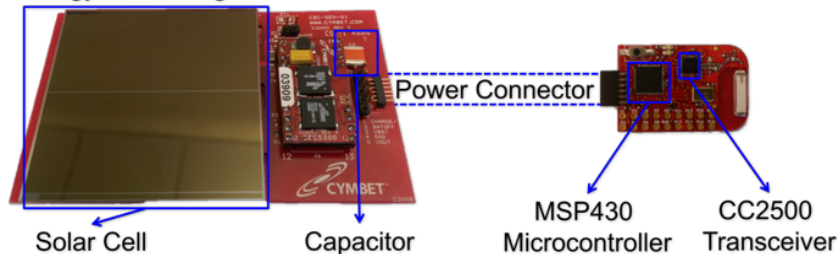
- Latency drops with increased σ
- The 99th percentile latency is better than that of Searchlight



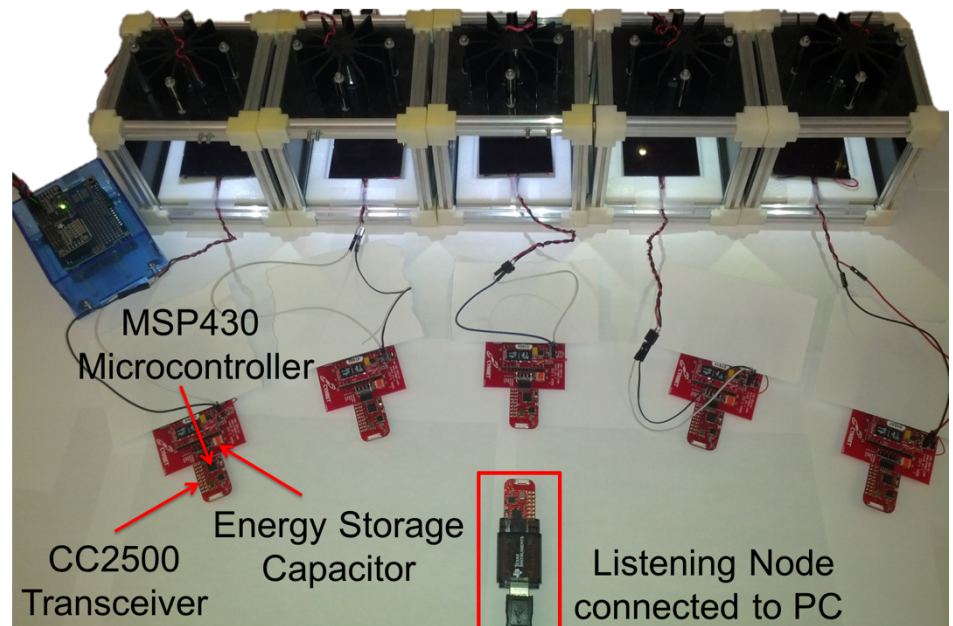
Experimental Evaluation

- TI eZ430-RF2500-SEH energy harvesting nodes
- Listen power: 67.08 mW, Transmit power: 56.29 mW
- Power budget: 1 mW, 5 mW
- Homogeneous nodes placed in proximity, $N = 5, 10$; $\sigma = 0.25, 0.5$

Energy Harvesting Power Source

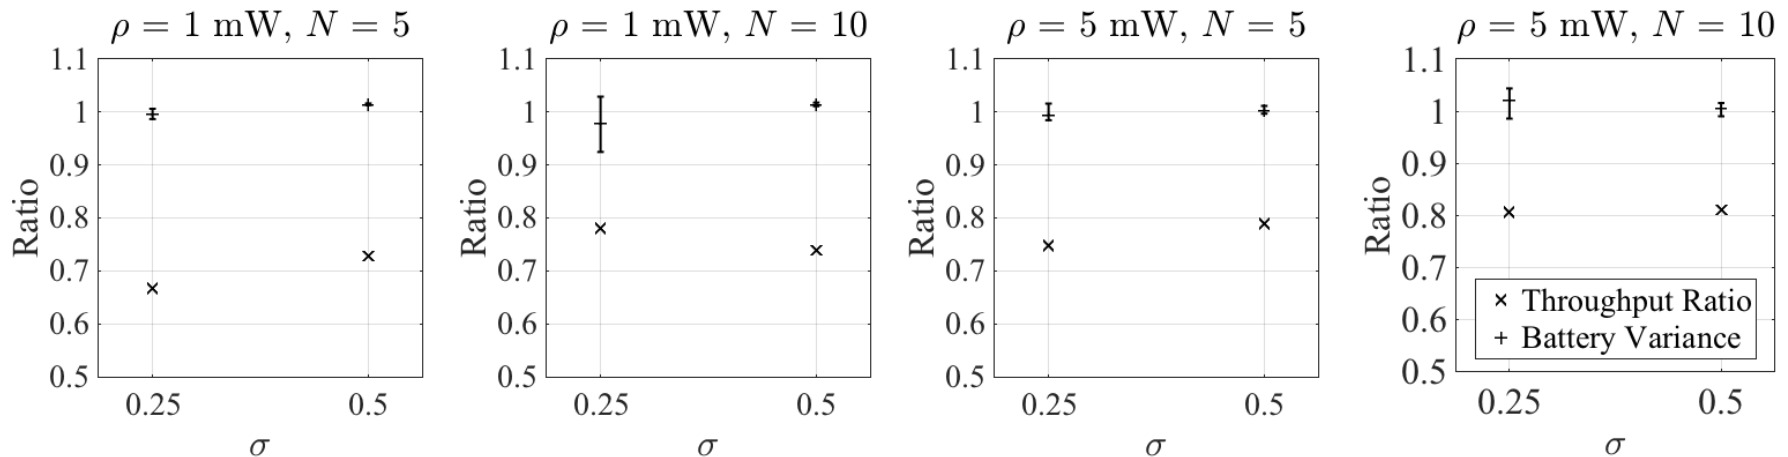


Light Control System + Solar Cells



Experimental Evaluation

- TI eZ430-RF2500-SEH energy harvesting nodes
- Listen power: 67.08 mW, Transmit power: 56.29 mW
- Power budget: 1 mW, 5 mW
- Homogeneous nodes placed in proximity, $N = 5, 10$; $\sigma = 0.25, 0.5$



- The experimental throughput is **67% – 81%** of the analytical throughput
- The experimental power consumption is within **3% – 7%** of the power budget
- In experiments, **EconCast** outperforms Panda by **8x – 11x** with $\sigma = 0.25$

Summary

- Broadcast throughput maximization among a set of heterogeneous energy-constrained ultra-low-power nodes
- ***EconCast*** – a simple asynchronous distributed protocol
- Prove that ***EconCast*** converges to the maximum achievable throughput
- Performance evaluation via simulations and experimentation
- Future directions:
 - Larger-scale network evaluation
 - Tradeoffs between throughput and latency
 - Integration with tracking applications

Thank you!

tingjun@ee.columbia.edu
www.columbia.edu/~tc2668/

Tingjun Chen, Javad Ghaderi, Dan Rubenstein, and Gil Zussman,
“Maximizing Broadcast Throughput Under Ultra-Low-Power Constraints”.

