

Maximizing Broadcast Throughput Under Ultra-Low-Power Constraints

Tingjun Chen[†], Javad Ghaderi[†], Dan Rubenstein[‡], Gil Zussman[†]
[†] Electrical Engineering, [‡] Computer Science, Columbia University
{tingjun@ee, jghaderi@ee, danr@cs, gil@ee}.columbia.edu

ABSTRACT

Wireless object tracking applications are gaining popularity and will soon utilize emerging ultra-low-power device-to-device communication. However, severe energy constraints require much more careful accounting of energy usage than what prior art provides. In particular, the available energy, the differing power consumption levels for listening, receiving, and transmitting, as well as the limited control bandwidth must all be considered. Therefore, we formulate the problem of maximizing the throughput among a set of heterogeneous broadcasting nodes with differing power consumption levels, each subject to a strict ultra-low-power budget. We obtain the oracle throughput (i.e., maximum throughput achieved by an oracle) and use Lagrangian methods to design EconCast – a simple asynchronous distributed protocol in which nodes transition between sleep, listen, and transmit states, and dynamically change the transition rates. We also show that EconCast approaches the oracle throughput. The performance is evaluated numerically and via extensive simulations and it is shown that EconCast outperforms prior art by 6x – 17x under realistic assumptions. Finally, we implement EconCast using the TI eZ430-RF2500-SEH energy harvesting nodes and experimentally show that in realistic environments it obtains 57% – 77% of the achievable throughput.

CCS Concepts

•Networks → Network protocol design; Network performance analysis; Wireless local area networks; •Mathematics of computing → Probabilistic algorithms;

Keywords

Ultra-low-power networking; Wireless networks; Sleep-wake protocols

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CoNEXT '16, December 12 - 15, 2016, Irvine, CA, USA

© 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-4292-6/16/12...\$15.00

DOI: <http://dx.doi.org/10.1145/2999572.2999608>

1. INTRODUCTION

Object tracking and monitoring applications are gaining popularity within the realm of Internet-of-Things [3]. One enabler of such applications is the growing class of ultra-low-power wireless nodes. An example is active tags that can be attached to physical objects, harvest energy from ambient sources, and communicate tag-to-tag toward gateways [20, 41]. Relying on node-to-node communications will require less infrastructure than traditional (RFID/reader-based) implementations. Therefore, as discussed in [7, 20, 34, 56], it is envisioned that such ultra-low-power nodes will facilitate tracking applications in healthcare, smart building, assisted living, manufacturing, supply chain management, and intelligent transportation.

A fundamental challenge in networks of ultra-low-power nodes is to schedule the nodes' sleep, listen/receive, and transmit events without coordination, such that they communicate effectively while adhering to their strict power budgets. For example, energy harvesting tags need to rely on the power that can be harvested from sources such as indoor-light or kinetic energy, which provide between 0.01 mW and 0.1 mW [21, 22] (for more details see the review in [51] and references therein). These power budgets are much lower than the power consumption levels of current low-power wireless technologies such as Bluetooth Low Energy [2] and ZigBee/802.15.4 [32] (usually at the order of 1 – 10 mW). On the other hand, Bluetooth Low Energy and ZigBee are designed to support data rates (up to a few Mbps) that are higher than required by the applications our work envisages supporting (less than a few Kbps).

In this paper, we formulate the problem of *maximizing broadcast throughput among energy-constrained nodes*. We design, analyze, and evaluate EconCast: **Energy-constrained BroadCast**. EconCast is an asynchronous distributed protocol in which nodes transition between sleep, listen/receive, and transmit states, while maintaining a power budget. The nodes and network we focus on have the following characteristics:

Broadcast: A transmission can be heard by all listening nodes in range.

Severe power constraints: The power budget is so limited that each node needs to spend most of its time in sleep state and the supported data rates can be of a few Kbps [22]. Traditional approaches that spend energy in order to improve

coordination (e.g., accurate clocks, slotting, synchronization) or form some sort of structure (e.g., routing tables and clusters) are too expensive given limited energy and bandwidth.

Unacquainted: Nodes do not require pre-existing knowledge of their environment (e.g., properties of neighboring nodes). This can result from the restricted power budget or from unanticipated environment changes due to altered energy sources and/or node mobility.

Heterogeneous: The power budgets and the power consumption levels can differ among the nodes.

Efficiently operating such structureless and ultra-low-power networks requires nodes to make their sleep, listen, or transmit decisions in a distributed manner. Therefore, we consider the fundamental problem of maximizing the rate at which the messages can be delivered (the actual content of the transmitted messages depends on the application). Namely, we focus on maximizing the broadcast throughput and consider two alternative definitions:

- **Groupput** – the total rate of successful bit transmissions to *all the receivers* over time. Groupput directly applies to tracking applications in which nodes utilize a neighbor discovery protocol to identify neighbors which are within wireless communication range [8, 27, 29, 42, 49, 53, 57, 60]. In such applications, broadcasting information to all other nodes in the network is important, allowing the nodes to transfer data more efficiently under the available power budgets. Groupput can also be applied to data flooding applications where the data needs to be collected at all the nodes in a network.
- **Anyput** – the total rate of successful bit transmissions to *at least one receiver* over time. It applies to delay-tolerant environments that utilize gossip-style methods to disseminate information. In traditional gossip communication, a node selects a communication partner in a deterministic or randomized manner. Then, it determines the content of the message to be sent based on a naive store-and-forward, compressive sensing [4, 23, 37, 38, 47, 48], or decentralized coding [15, 28]. As another example, in delay-tolerant applications, data transmission may get disrupted or lost due to the limits of wireless radio range, sparsity of mobile nodes, or limited energy resources, a node may wish to send its data to any available receiver.

First, we derive oracle throughput (i.e., maximum throughput achieved by an oracle). This is done both for the groupput and anyput and it is shown that the value can be efficiently computed. Then, we use Lagrangian methods and a Q-CSMA (Queue-based Carrier Sense Multiple Access) approach to design EconCast. The protocol has variants that maximize both groupput and anyput. In both of them, nodes dynamically adapt their transition rates between sleep, listen, and transmit states based on (i) the energy available at the node and (ii) the number (or existence) of other active listeners. To support the latter, a listening node emits a low-cost informationless “ping” which can be picked up by other listening nodes, allowing them to estimate the number (or existence) of active listeners. We briefly discuss how this

method helps increasing the throughput and the implementation aspects. We analyze the performance of EconCast and prove that, in theory, it converges to the oracle throughput.

We evaluate EconCast numerically and via extensive simulations under a wide range of power budgets, and listen/transmit power consumption levels, and for various heterogeneous nodes. Specifically, numerical results show that EconCast outperforms prior art (Panda [42], Birthday [43], and Searchlight [5]) by a factor of 6x – 17x under realistic assumptions. In addition to throughput, we consider the performance in terms of burstiness and latency.

We implement EconCast using the TI eZ430-RF2500-SEH energy harvesting nodes and experimentally show that in practice it obtains 57% – 77% of the achievable throughput. Moreover, we compare the throughput obtained experimentally to analytical results for Panda [42] (where the analytical results are usually better than the experimental performance) and show that, for example, EconCast outperforms Panda by 8x – 11x.

We note that the design of EconCast does not assume a specific topology (nodes do not know anything about their neighbors). Yet, in this paper, we mainly focus on a clique topology (i.e., nodes are within the communication range of each other), since it lends itself to analysis. We briefly extend the analytic results to non-clique topologies and also evaluate the performance for such networks.

To summarize, the main contributions of this paper are: (i) a distributed asynchronous protocol for a heterogeneous collection of energy-constrained wireless nodes, that can obtain throughput that approaches the maximum possible, (ii) efficient methods to compute the oracle throughput, (iii) extensive performance evaluation of the protocol.

The rest of the paper is organized as follows. We discuss related work in Section 2 and formulate the problem in Section 3. In Section 4, we present methods to compute the oracle throughput. We present EconCast in Section 5 and outline the proof of the main theoretical result in Section 6. We evaluate EconCast numerically and via simulations and compare to related work in Section 7. In Section 8, we discuss the experimental implementation and evaluation. We conclude in Section 9. Due to space constraints, some of the proofs are omitted and can be found in the technical report [12].

2. RELATED WORK

There is vast amount of related literature in sensor networking and neighbor discovery that tries to limit energy consumption. Within this large body of work, most of the protocols do not explicitly account for different listen and transmit power consumption levels of the nodes [5, 6, 8, 9, 16, 26, 43, 46, 49, 50, 52–54, 59, 60], or do not account for different power budgets [14, 17, 52–54, 59]. They mostly use a duty cycle during which nodes sleep to conserve energy and when nodes are simultaneously awake, a pre-determined listen-transmit sequence with an unalterable power consumption level is used. However, for ultra-low-power nodes constrained by severe power budgets, the appropriate amount

\mathcal{N}, N	Set of nodes, number of nodes
L_i, \mathbf{L}	Node i 's listen power consumption (W), $\mathbf{L} = [L_i]$
X_i, \mathbf{X}	Node i 's transmit power consumption (W), $\mathbf{X} = [X_i]$
$\rho_i, \boldsymbol{\rho}$	Node i 's power budget (W), $\boldsymbol{\rho} = [\rho_i]$
b_i	Energy storage level of node i (J)
\mathbf{w}, \mathcal{W}	Network state, the set of collision-free states
$\alpha_i, \boldsymbol{\alpha}$	Fraction of time node i listens, $\boldsymbol{\alpha} = [\alpha_i]$
$\beta_i, \boldsymbol{\beta}$	Fraction of time node i transmits, $\boldsymbol{\beta} = [\beta_i]$
$\gamma, \hat{\gamma}$	Indicator if existing some nodes listening, its estimated value
c, \hat{c}	Number of nodes listening, its estimated value
ν	Indicator if there is exactly one node transmitting
$\pi_{\mathbf{w}}, \boldsymbol{\pi}$	Fraction of time the network is in $w \in \mathcal{W}$, $\boldsymbol{\pi} = [\pi_{\mathbf{w}}]$
$\mathcal{T}_g, \mathcal{T}_a$	Groupup and Anyput of the network
$\mathcal{T}_{\mathbf{w}}$	Throughput of state $\mathbf{w} \in \mathcal{W}$
\mathcal{T}^*	Oracle throughput
$\eta_i, \boldsymbol{\eta}$	Lagrangian multiplier of node i , $\boldsymbol{\eta} = [\eta_i]$

Table 1: Nomenclature

of time a node sleeps should explicitly depend on the relative listen and transmit power consumption levels. These prior approaches achieve throughput levels which are much below optimal (and hence much below what EconCast can achieve). Additionally, there are protocols that often require some explicit coordination (e.g., slotting [14, 43, 46, 59], or explicitly require exchange of parameters [42], which are not suitable for emerging ultra-low-power nodes.

From the theoretical point of view, our approach is inspired by the prior work on network utility maximization (e.g., [13, 30, 35, 36]), and queue-based CSMA literature (e.g., [18, 19, 25, 33, 40, 58]). However, the problem considered in this paper is not a simple extension of the prior work for two reasons. First, in the past work on CSMA and network utility maximization, nodes or links make decisions based on the relative sizes of queues. Often, a queue is a backlog of data to send or the available energy. Prior work that considers the latter (e.g., [10, 31, 39]) uses the energy only for transmission, while listening is “free”, which is a very different paradigm than the one considered in this paper. Second, in our setting, the queue “backlogs” energy but there is no clear mapping as previously assumed from energy to successful transmission. A node’s listen or transmit events will relieve the backlog, but *do not* increase utility (throughput) unless other nodes are appropriately configured (i.e., transmitting when no listening nodes exist or listening when no transmitting nodes exist does not increase the throughput). This coordination of state among nodes to utilize their energy makes the considered problem more challenging.

Finally, we note that our approach should be amenable to emerging physical layer broadcast methods such as backscatter [34, 44].

3. MODEL AND PROBLEM FORMULATION

We consider a network of N energy-constrained nodes whose objective is to distributedly maximize the broadcast throughput among them. The set of nodes is denoted by \mathcal{N} . Table 1 summarizes the notations.

3.1 Basic Node Model

Power consumption: A node $i \in \mathcal{N}$ can be in one of three

states: *sleep* (s), *listen/receive*¹ (l), and *transmit* (x), and the respective power consumption values are 0, L_i (W), and X_i (W).² These power consumption levels are based on hardware characteristics.

Power budget: Each node i has a *power budget* of ρ_i (W). This budget can be the rate at which energy is harvested by an energy harvesting node or a limit on the energy spending rate such that the node can maintain a certain lifetime. In practice, the power budget may vary with time [21, 22] and the distributed protocol should be able to adapt. For simplicity, we assume that the power budget is constant with respect to time. However, the analysis can be easily extended to the case with time-varying power budget with the same constant mean. Each node i also has an energy storage (e.g., a battery or a capacitor) whose level at time t is denoted by $b_i(t)$.

Severe Power Constraints: Intermittently connected energy-constrained nodes cannot rely on complicated synchronization or structured routing approaches.

Unacquainted: Low bandwidth implies that each node i must operate with very limited (i.e., no) knowledge regarding its neighbors, and hence, does not know or use the information (ρ_j, L_j, X_j) of the other nodes $j \neq i$.

3.2 Architecture Assumptions

We assume that there is only *one frequency channel* and *a single transmission rate* is used by all nodes in the transmit state. Similar to CSMA, nodes perform *carrier sensing* prior to attempting transmission to check the availability of the medium. Energy-constrained nodes can only be awake for very short periods, and therefore, the likelihood of overlapping transmissions is negligible.

We also assume that *a node in the listen state can send out low-cost, informationless “pings”* which can be picked up by other listening nodes, allowing them to estimate the number (or existence) of active listeners. We explain in Section 5 how this property will help us develop a distributed protocol and in Section 8, we provide practical means by which such estimates can be obtained.

3.3 Model Simplifications

At any time t , the network state can be described as a vector $\mathbf{w}(t) = [w_i(t)]$, where $w_i(t) \in \{s, l, x\}$ represents the state of node i . While the distributed protocol EconCast (described in Section 5) can operate in general scenarios, for analytical tractability, we make the following assumptions:

- The network is a *clique*.³
- Nodes can perform *perfect carrier sensing* in which the propagation delay is assumed to be zero.

These assumptions are suitable in the envisioned applications where the distances between nodes are small. Under these assumptions, the network states can be restricted to

¹We refer the *listen* and *receive* states synonymously as the power consumption in both states is similar.

²The actual power consumption in the *sleep* state, which may be non-zero, can be incorporated by reducing ρ_i , or increasing both L_i and X_i , by the sleep power consumption.

³We also investigate non-clique networks in Section 4.3.

the set of *collision-free* states, denoted by \mathcal{W} (i.e., states in which there is *at most* one node in transmit state). This reduces the size of the state space from 3^N to $(N+2)2^{N-1}$.

Let $\gamma_{\mathbf{w}} \in \{0, 1\}$ indicate whether there *exists* some nodes listening in state \mathbf{w} and let $c_{\mathbf{w}}$ be the number of listeners in state \mathbf{w} . We use $\nu_{\mathbf{w}} \in \{0, 1\}$ as an indicator which is equal to 1 if there is *exactly* one transmitter in state \mathbf{w} and is 0 otherwise. Based on these indicator functions, two measures of broadcast throughput, *groupput* and *anyput*, and the throughput of a given network state \mathbf{w} are defined below.

DEFINITION 1 (GROUPPUT). *The groupput, denoted by \mathcal{T}_g , is the aggregate throughput of the transmissions received by all the receivers, where each transmitted bit is counted once per receiver to which it is delivered, i.e.,*

$$\mathcal{T}_g = \lim_{T \rightarrow \infty} \frac{1}{T} \int_{t=0}^T \nu_{\mathbf{w}(t)} c_{\mathbf{w}(t)} dt. \quad (1)$$

DEFINITION 2 (ANYPUT). *The anyput, denoted by \mathcal{T}_a , is the aggregate throughput of the transmissions that are received by at least one receiver, i.e.,*

$$\mathcal{T}_a = \lim_{T \rightarrow \infty} \frac{1}{T} \int_{t=0}^T \nu_{\mathbf{w}(t)} \gamma_{\mathbf{w}(t)} dt. \quad (2)$$

DEFINITION 3 (NETWORK STATE THROUGHPUT). *The throughput associated with a given network state $\mathbf{w} \in \mathcal{W}$, denoted by $\mathcal{T}_{\mathbf{w}}$, is defined as*

$$\mathcal{T}_{\mathbf{w}} = \begin{cases} \nu_{\mathbf{w}} c_{\mathbf{w}}, & \text{for Groupput} \\ \nu_{\mathbf{w}} \gamma_{\mathbf{w}}, & \text{for Anyput} \end{cases} \quad (3)$$

Note that without energy constraints, the oracle (maximum) groupput is $(N-1)$ and is achieved when some node always transmits and the remaining $(N-1)$ nodes always listen and receive the transmission. Similarly, the oracle (maximum) anyput without energy constraints is 1 and is achieved when some node always transmits and some other node always listens and receives the transmission.

3.4 Problem Formulation

Define $\pi_{\mathbf{z}}$ as the fraction of time the network spends in a given state $\mathbf{z} \in \mathcal{W}$, i.e.,

$$\pi_{\mathbf{z}} = \lim_{T \rightarrow \infty} \frac{1}{T} \int_{t=0}^T \mathbf{1}_{\{\mathbf{w}(t)=\mathbf{z}\}} dt, \quad (4)$$

where $\mathbf{1}_{\{\mathbf{w}(t)=\mathbf{z}\}}$ is the indicator function which is 1, if the network is with state \mathbf{z} at time t , and is 0 otherwise. Correspondingly, denote $\boldsymbol{\pi} = [\pi_{\mathbf{w}}]$.

Below, we define the *energy-constrained throughput maximization problem (P1)* where the fractions of time each node spends in sleep, listen, and transmit states are assigned while the node maintains the power budget. Define variables $\alpha_i, \beta_i \in [0, 1]$ as the fraction of time node i spends in listen and transmit states, respectively. The fraction of time it spends in sleep state is simply $(1 - \alpha_i - \beta_i)$. In view of (1)

– (4), (P1) is given by

$$(P1) \quad \max_{\boldsymbol{\pi}} \quad \sum_{\mathbf{w} \in \mathcal{W}} \pi_{\mathbf{w}} \mathcal{T}_{\mathbf{w}} \quad (5)$$

$$\text{subject to} \quad \alpha_i L_i + \beta_i X_i \leq \rho_i, \quad \forall i \in \mathcal{N}, \quad (6)$$

$$\alpha_i = \sum_{\mathbf{w} \in \mathcal{W}_i^l} \pi_{\mathbf{w}}, \quad \beta_i = \sum_{\mathbf{w} \in \mathcal{W}_i^x} \pi_{\mathbf{w}}, \quad (7)$$

$$\sum_{\mathbf{w} \in \mathcal{W}} \pi_{\mathbf{w}} = 1, \quad \pi_{\mathbf{w}} \geq 0, \quad \forall \mathbf{w} \in \mathcal{W}, \quad (8)$$

where \mathcal{W}_i^l and \mathcal{W}_i^x are the sets of states $\mathbf{w} \in \mathcal{W}$ in which $w_i = l$ and $w_i = x$, respectively. Each node is constrained by a power budget, as described in (6), and (8) represents the fact that at any time, the network operates in one of the collision-free states $\mathbf{w} \in \mathcal{W}$.

Based on the solution to (P1), the maximum throughput is achievable by an *oracle* that can schedule nodes' sleep, listen, and transmit periods, in a centralized manner. Therefore, we define the maximum value obtained by solving (P1) as the *oracle throughput*, denoted by \mathcal{T}^* . Respectively, we define the *oracle groupput* and *oracle anyput* as \mathcal{T}_g^* and \mathcal{T}_a^* .

To evaluate EconCast, it is essential to compare its performance to the oracle throughput. However, (P1) is a Linear Program (LP) over an exponentially large number of variables (i.e., $|\mathcal{W}|$ is exponential in N) and is computationally expensive to solve. In Section 4, we show how to convert (P1) to another optimization problem with only a linear number of variables. Note that the solution to (P1) only provides the optimal fraction of time each node should spend in sleep, listen, and transmit states, but *does not* indicate how the nodes can make their individual sleep, listen, and transmit decisions locally. Therefore, in Section 5, we focus on the design of EconCast that makes these decisions based on (P1).

4. ORACLE THROUGHPUT

In this section, we present an equivalent LP formulation for (P1) in a clique network which only has a linear number of variables. We also derive both an upper and a lower bound for the oracle groupput in non-clique topologies which will be used later for evaluating the performance of EconCast in non-clique topologies.

Recall that α_i and β_i are the fraction of time node i spends in listen and transmit states, respectively. We can rewrite the constraints in (P1) as follows

$$\alpha_i L_i + \beta_i X_i \leq \rho_i, \quad \forall i \in \mathcal{N}, \quad (9)$$

$$\alpha_i + \beta_i \leq 1, \quad \forall i \in \mathcal{N}, \quad (10)$$

$$\sum_{i \in \mathcal{N}} \beta_i \leq 1. \quad (11)$$

Specifically, (9) is the usual power budget constraint on each node $i \in \mathcal{N}$, and (10) is due to the fact that a node can only operate in one state at any time. We remark that energy-constrained nodes can only be awake for very small fractions of time (i.e., $\alpha_i + \beta_i \ll 1$), and therefore (10) may be redundant. Finally, collision-free operation in a clique network when *at most* one transmitter can be present at any time imposes (11) which bounds the sum of the transmit fractions by 1.

4.1 Oracle Groupput in a Clique

To maximize the groupput (1), it suffices that any node only listens when there is another transmitter, since listening when no one transmits wastes energy. Namely, the fraction of time node i listens cannot exceed the aggregate fraction of time all other nodes transmit, i.e.,

$$\alpha_i \leq \sum_{j \neq i} \beta_j, \forall i \in \mathcal{N}. \quad (12)$$

Since a node only listens when there exists *exactly one* transmitter, every listen counts as a reception, and the groupput of a node (i.e., the throughput it receives from all other nodes) is simply the fraction of time it spends in listen state α_i . Therefore, the groupput in a clique network simplifies to $\sum_{i \in \mathcal{N}} \alpha_i$. The *oracle groupput*, denoted by \mathcal{T}_g^* , can be obtained by solving the following maximization problem

$$\begin{aligned} \text{(P2)} \quad \mathcal{T}_g^* := & \max_{\alpha, \beta} \sum_{i \in \mathcal{N}} \alpha_i \\ & \text{subject to} \quad (9) - (12). \end{aligned} \quad (13)$$

(P2) is an LP consisting of $2N$ variables and $(3N + 1)$ constraints (i.e., solving for α and β given inputs of N , ρ , \mathbf{L} , and \mathbf{X}). On a conventional laptop running Matlab, this computation for thousands of nodes takes seconds. Moreover, we show that the oracle groupput obtained by solving (P2) is indeed *achievable* by an oracle which can schedule nodes' listen and transmit periods. This result is summarized in the following lemma and the proof is in [12].

LEMMA 1. *The (rational-valued) solution (α^*, β^*) to (P2) can be feasibly scheduled by an oracle in a fixed-size slotted environment via a periodic schedule, (perhaps) after a one-time energy accumulation interval.*

For the case of *homogeneous* nodes (i.e., $\rho_i = \rho$, $L_i = L$, $X_i = X$, $\forall i \in \mathcal{N}$) where nodes are sufficiently energy-constrained (i.e., (9) dominates (10)), the closed-form solution to (P2) (note that the equalities hold for equations (9) and (12)⁴) is given by

$$\beta^* = \rho / (X + (N - 1)L), \alpha^* = (N - 1)\beta^*, \mathcal{T}_g^* = N\alpha^*.$$

4.2 Oracle Anyput in a Clique

The oracle anyput is obtained based on the observation that a transmission only occurs when there is *at least* one listener. We define additional variables $\chi_{i,j}$ as the fraction of time node j receives a transmission from node i , for the following two constraints

$$\beta_i \leq \sum_{j \neq i} \chi_{i,j}, \quad \forall i \in \mathcal{N}, \quad (14)$$

$$\alpha_j = \sum_{i \neq j} \chi_{i,j}, \quad \forall j \in \mathcal{N}. \quad (15)$$

The *oracle anyput*, denoted by \mathcal{T}_a^* , can be obtained by solving the following maximization problem

$$\begin{aligned} \text{(P3)} \quad \mathcal{T}_a^* := & \max_{\alpha, \beta} \sum_{i \in \mathcal{N}} \beta_i \\ & \text{subject to} \quad (9) - (11), (14), \text{ and } (15). \end{aligned} \quad (16)$$

⁴This can be proved by contradiction. The details can be found in [12].

First, (14) ensures that when node i transmits, there is always *at least* one other node than can receive this transmission. Then, (15) makes sure that in the optimal schedule, the fraction of time node j listens is large enough to cover all the transmissions it receives. Therefore, (P3) maximizes the anyput by ensuring that every transmission is received by *at least* one node.

For the case of homogeneous nodes, the closed-form solution to (P3) is given by

$$\beta^* = \alpha^* = \rho / (X + L), \mathcal{T}_a^* = N\beta^*.$$

4.3 Oracle Groupput in Non-cliques

The problem formulations (P1) – (P3) so far have assumed a clique network. Obtaining the exact maximum groupput for non-cliques (denoted by \mathcal{T}_{nc}^*) is difficult. This is because a node may receive simultaneous transmissions from two nodes which are not within communication range of each other. As explained before, listen and transmit events are rare within energy-constrained nodes. Therefore, the likelihood of simultaneous transmissions is small and it is expected to have minimal impact on the throughput.

We present both an upper bound $\overline{\mathcal{T}}_{nc}^*$ and a lower bound $\underline{\mathcal{T}}_{nc}^*$ on the maximum groupput in non-clique topologies. In the scenarios where $\overline{\mathcal{T}}_{nc}^*$ and $\underline{\mathcal{T}}_{nc}^*$ are the same, the exact maximum groupput \mathcal{T}_{nc}^* can be obtained. The lower bound $\underline{\mathcal{T}}_{nc}^*$ is obtained by solving (P2) but replace constraint (12) by

$$\alpha_i \leq \sum_{i \in \mathcal{N}(i)} \beta_j, \quad \forall i \in \mathcal{N},$$

where $\mathcal{N}(i)$ is the set of neighboring nodes of node i . This ensures that the fraction of time node i listens cannot exceed the sum of its neighboring nodes' fractions of transmissions. The upper bound $\overline{\mathcal{T}}_{nc}^*$ is obtained by solving (P2) in which the constraint (11) is removed. This allows overlapping transmissions which can possibly happen in non-cliques. Numerical results show that with certain topologies, $\overline{\mathcal{T}}_{nc}^* = \underline{\mathcal{T}}_{nc}^*$ holds, resulting in the exact maximum groupput \mathcal{T}_{nc}^* . In Section 7.5, we compute \mathcal{T}_{nc}^* and evaluate the performance of EconCast in such scenarios.

5. DISTRIBUTED PROTOCOL

In this section, we describe EconCast from the perspective of a single node that transitions between sleep, listen, and transmit states, under a power budget. Since we focus on a single node i , in parts of this section, we drop the subscript i of previously defined variables for notational compactness.

5.1 A Simple Heterogeneous Example

To better understand the challenges faced in designing EconCast, consider a simple example of 4 nodes, all having identical listen and transmit costs $L_i = X_i = 1$ mW ($i = 1, 2, 3, 4$), but different power budgets ρ_i , as indicated in Table 2. Table 2 also shows the percentage of time each node spends in listen and transmit states (α_i^*, β_i^*) ($i = 1, 2, 3, 4$) such that the groupput is maximized by solving (P1). It also shows the percentage of time each node spends in transmit state *when awake* (i.e., $\frac{100 \cdot \beta_i^*}{\alpha_i^* + \beta_i^*} \%$).

Node	1	2	3	4
Power Budget: ρ_i (mW)	0.005	0.01	0.05	0.1
Awake(%): $\alpha_i^* + \beta_i^*$	0.5	1.0	5.0	10.0
Transmit when Awake(%)	20.0	22	53.6	65.7

Table 2: A simple example in a heterogeneous network.




Figure 1: The node’s states and transition rates.

If, instead, all nodes have the same power budget of $\rho_i = 0.1$ mW, the percentage of time each node spends in transmit state when awake is 25% (with $\alpha_i^* = 0.075$, $\beta_i^* = 0.025$, $i = 1, 2, 3, 4$). Note that in the above example, the power budget of node 4 remains unchanged but changes in other nodes’ power budgets shift the percentage of time it should transmit when awake from 25% to 65.7%. This clearly shows that the partitioning of a node’s power budget among listen and transmit states is highly dependent on other nodes’ properties. *However, we will show that if a node does not know the properties of its neighbors, an optimal configuration can be obtained without explicitly solving (P1).*

5.2 Protocol Description

To clearly present EconCast, we start from a theoretical framework and slowly build on it to address practicalities. As mentioned in Section 3, a node can be in one of three states: sleep (s), listen (l), and transmit (x). As depicted in Figure 1, it must pass through the listen state to transition between sleep and transmit states. The time duration a node spends in a given state u before transitioning to state v is exponentially distributed with rate $\lambda_{uv}(t)$. These *transition rates* can be adjusted over time. We remark that sending packets with exponentially distributed length (i.e., a node transitions from transmit state to listen state with a rate λ_{xl}) is impractical. However, it can be shown that this is equivalent to continuously transmitting back-to-back unit-length packets with probability $(1 - \lambda_{xl})$ if $\lambda_{xl} \in [0, 1]$, which is indeed the case in EconCast.

The throughput (5) as a function of π_w is controlled by appropriately adjusting the transition rates between different states of each node. EconCast determines in a distributed manner how these adjustments are performed over time. Roughly speaking, each node adjusts its transition rates $\lambda_{uv}(t)$ based on limited information that can be obtained in practice, which includes

- Its power consumption levels, L and X , and energy storage level $b(t)$.
- A sensing of transmit activity of other nodes over the channel (CSMA-like carrier sensing).
- A count of other active listeners (for groupput maximization), $c(t)$, or an indicator of whether there are any active listeners (for anyput maximization), $\gamma(t)$. In practice, $c(t)$

and $\gamma(t)$ may not be accurate, and we denote $\hat{c}(t)$ and $\hat{\gamma}(t)$ as their estimated values.

We note that in EconCast, unlike in previous work such as Panda [42], each node *does not* need to know the number of nodes in the network, N , and the power budgets and power consumption levels of other nodes. Furthermore, a node does not need to know its power budget ρ explicitly (e.g., in the case of energy harvesting [41]), although this knowledge can be incorporated, if available.

Under EconCast, a node sets $\lambda_{sl}(t)$ as an increasing function of the available stored energy, $b(t)$, to more aggressively exit sleep state. Furthermore, it sets $\lambda_{lx}(t)$ as an increasing function of the number of listeners, $\hat{c}(t)$, to enter transmit state more frequently when more nodes are listening. We will describe how these functions are chosen in Section 5.5.

5.3 Estimating Active Listeners: Pings

As described above, an important input to EconCast is the number of active listeners $c(t)$ (for groupput) or the indicator of existence of active listeners $\gamma(t)$ (for anyput). We now discuss the estimation of $\hat{c}(t)$ or $\hat{\gamma}(t)$. Recall from Section 3 that nodes can send out periodic *pings* that any other listener can receive. The pings need not carry any explicit information and are potentially significantly cheaper and shorter than control packet transmissions (e.g., an ACK). Therefore, they consume less power and take much less time than a minimal data transmission.

Consider the case in which all nodes are required to send pings at a pre-determined rate and the power consumption is accounted for in the listening power consumption L . In such a case, a fellow listener detecting such pings (e.g., using a simple energy detector) can use the count of such pings in a given period of time, or the inter-arrival times of pings, to estimate the number of active listeners $c(t)$. Estimating $\gamma(t)$ is even easier by detecting the existence of any ping. In general, the estimates do not need to be accurate for EconCast to function, although poor estimates are expected to reduce throughput.

5.4 Two Variants of EconCast

We now address the incorporation of the estimates $\hat{c}(t)$ and $\hat{\gamma}(t)$ into EconCast. We present two versions of EconCast which only differ when a node is in transmit state:

- EconCast-C (the capture version): a node may “capture” the channel and transmit for an exponential amount of time (i.e., several back-to-back packets). When each packet transmission is completed, the transmitter listens for pings for a fixed-length *pinging interval*. Each successful recipient of the transmission initiate one ping at time chosen uniformly at random on this interval. The transmitter then estimates $\hat{c}(t)$ or $\hat{\gamma}(t)$ based on the count of pings received and adjusts $\lambda_{xl}(t)$ (as described in Section 5.5). In Section 8.3, we discuss the experimental implementation of this process.
- EconCast-NC (the non-capture version): a node always releases the channel after one packet transmission. Each

node continuously pings and receives pings from other nodes when listening, estimates $\hat{c}(t)$ or $\hat{\gamma}(t)$, and adjusts $\lambda_{lx}(t)$ (as described in Section 5.5).

EconCast-C is significantly easier to implement since the estimates are only needed for the transmitter right after each packet transmission. The probability that the same transmitter will continue transmitting depends on the estimates $\hat{c}(t)$ or $\hat{\gamma}(t)$. Therefore, our implementation and experimental evaluations in Section 8 focus on EconCast-C.

5.5 Setting Transition Rates

Consider a node running EconCast. Time is broken into intervals of length τ_k ($k = 1, 2, \dots$). The k -th interval is from time t_{k-1} to time t_k and we let $t_0 = 0$. EconCast takes input of two internal variables:

- η is a multiplier which is updated at the beginning of each time interval. Let $b[k]$ ($k = 0, 1, \dots$) denote the energy storage level at the end of the k -th time interval. Let $(\cdot)^+$ denote $\max(0, \cdot)$ and $\eta[k]$ is updated as follows

$$\eta[k] = \left(\eta[k-1] - \frac{\delta_k}{\tau_k} \cdot (b[k] - b[k-1]) \right)^+, \quad (17)$$

in which $\delta_k \in (0, 1)$ is a step size and $b[k] = b(t_k)$. We use square brackets here to imply that the multiplier $\eta[k]$ remains constant for $t \in [t_k, t_{k+1})$.

- $A(t)$ is the carrier sensing indicator of a node, which is 1 when the node does not sense any ongoing transmission, and is 0 otherwise. Carrier sensing forces a node to “stick” to its current state. When receiving an ongoing transmission, a node in listen state will not exit the listen state until it finishes receiving the full transmission, and a node in sleep state will not leave the sleep state (i.e., it enters the listen state but immediately leaves when it hears the ongoing transmission by performing carrier sensing).

The transition rates are described as follows (the superscripts C and N denote EconCast-C and EconCast-NC). For groupput maximization, at any time t in the k -th interval,

$$\lambda_{sl}(t) = A(t) \cdot \exp[-\eta[k]L/\sigma], \quad (18a)$$

$$\lambda_{ls}(t) = A(t), \quad (18b)$$

$$\lambda_{lx}^C(t) = A(t) \cdot \exp[\eta[k](L - X)/\sigma], \quad (18c)$$

$$\lambda_{lx}^N(t) = A(t) \cdot \exp[\eta[k](L - X)/\sigma + \hat{c}(t)/\sigma], \quad (18d)$$

$$\lambda_{xl}^C(t) = \exp[-\hat{c}(t)/\sigma], \quad (18e)$$

$$\lambda_{xl}^N(t) = 1. \quad (18f)$$

For anyput maximization, $\hat{c}(t)$ is replaced with $\hat{\gamma}(t)$. Theorem 1 below states the main result of this paper. We outline the proof and explain the intuition behind the protocol in Section 6.

THEOREM 1. *Let $\sigma \rightarrow 0$ and select parameters δ_k and τ_k properly (e.g., $\delta_k = 1/[(k+1)\log(k+1)]$ and $\tau_k = k$). Under perfect knowledge of $c(t)$ or $\gamma(t)$, the average throughput of EconCast (\mathcal{T}_g or \mathcal{T}_a) converges to the oracle throughput (\mathcal{T}_g^* or \mathcal{T}_a^*) given by (P1).*

5.6 Stability and Choice of σ , δ_k , and τ_k

EconCast is adaptive and, as expected, it must deal with the tradeoff of “adapting quickly but poorly” to “adapting optimally but slowly”. This adaptation manifests itself into the parameters σ , δ_k , and τ_k . When σ is increased, the throughput is less bursty (nodes transition from transmit state to listen state more frequently). However, the resulting throughput also decreases with respect to increased σ , as we will describe in Section 6.

Under a given value of σ , each node continuously adjusts the rates $\lambda_{uv}(t)$ based on its multiplier η according to (17), which is a function of the ratio δ_k/τ_k . Small δ_k/τ_k ratios make smaller changes of η over time, and lead to longer convergence time to the “right” multiplier values. In contrast, larger δ_k/τ_k ratios make η oscillate more wildly near the optimal value, such that the performance of EconCast is further from the optimal. Although the guaranteed convergence requires careful choices of the parameters (as stated in Theorem 1), in practice, we can choose $\delta_k = \delta$ and $\tau_k = \tau$ for some small constant δ and large constant τ .

6. PROOF OUTLINE OF THEOREM 1

In this section we provide an outline of the proof of Theorem 1. The complete proof can be found in [12]. The proof is based on a Markov Chain Monte Carlo (MCMC) approach [18, 25, 40] from statistical physics.

First, note that if the vector of multipliers $\boldsymbol{\eta} = [\eta_i]$ freezes, EconCast generates the network state distribution described in the following lemma.

LEMMA 2. *With fixed $\boldsymbol{\eta}$, the network Markov chain, resulted from overall interactions among the nodes according to the transition rates (18), has the steady state distribution*

$$\pi_{\mathbf{w}}^{\boldsymbol{\eta}} = \frac{1}{Z^{\boldsymbol{\eta}}} \exp \left[\frac{1}{\sigma} \left(\mathcal{T}_{\mathbf{w}} - \sum_{i:w_i=l} \eta_i L_i - \sum_{i:w_i=x} \eta_i X_i \right) \right], \quad (19)$$

where $Z^{\boldsymbol{\eta}}$ is a normalizing constant so that $\sum_{\mathbf{w} \in \mathcal{W}} \pi_{\mathbf{w}}^{\boldsymbol{\eta}} = 1$.

PROOF. The proof is followed by checking that the steady state distribution (19) satisfies the detailed balance equations of the network Markov chain. Details can be found in [12]. \square

We then present an optimization problem (P4) as follows

$$\begin{aligned} \text{(P4)} \quad & \max_{\boldsymbol{\pi}} \sum_{\mathbf{w} \in \mathcal{W}} \pi_{\mathbf{w}} \mathcal{T}_{\mathbf{w}} - \sigma \sum_{\mathbf{w} \in \mathcal{W}} \pi_{\mathbf{w}} \log \pi_{\mathbf{w}} \quad (20) \\ & \text{subject to} \quad (6), (7), \text{ and } (8), \end{aligned}$$

where σ is the positive constant used in EconCast (the counterpart in statistical physics is the temperature in systems of interacting particles). Note that (P4) is a concave maximization problem and as $\sigma \rightarrow 0$, the optimal value of (P4) approaches that of (P1). To solve (P4), consider the Lagrangian function $\mathcal{L}(\boldsymbol{\pi}, \boldsymbol{\eta})$ formulated by moving the energy constraint (6) into the objective (20) with a Lagrange multiplier $\eta_i \geq 0$ for each node i , i.e.,

$$\begin{aligned} \mathcal{L}(\boldsymbol{\pi}, \boldsymbol{\eta}) = & \sum_{\mathbf{w} \in \mathcal{W}} \pi_{\mathbf{w}} \mathcal{T}_{\mathbf{w}} - \sigma \sum_{\mathbf{w} \in \mathcal{W}} \pi_{\mathbf{w}} \log \pi_{\mathbf{w}} \\ & - \sum_{i \in \mathcal{N}} [\eta_i (\alpha_i L_i + \beta_i X_i - \rho_i)]. \quad (21) \end{aligned}$$

Algorithm 1 Gradient Descent Algorithm

Input parameters: σ, ρ, \mathbf{L} , and \mathbf{X}
Initialization: $\alpha_i(0) = \beta_i(0) = \eta_i(0) = 0, \forall i \in \mathcal{N}$
1: **for** $k = 1, 2, \dots$ **do**
2: $\delta(k) = 1/k$, compute $\boldsymbol{\pi}(k)$ from (19) using $\boldsymbol{\eta} = \boldsymbol{\eta}(k)$
3: **for** $i = 1, 2, \dots, N$ **do**
4: Update $\eta_i(k), \alpha_i(k)$, and $\beta_i(k)$ according to (23), (24)

In view of (7) and (8), given a vector of multipliers $\boldsymbol{\eta}$, it can be shown that the optimal $\boldsymbol{\pi}^\eta = [\pi_{\mathbf{w}}^\eta]$ that maximizes $\mathcal{L}(\boldsymbol{\pi}, \boldsymbol{\eta})$ is exactly given by (19). Thus if EconCast knows the optimal choice of Lagrange multipliers, it can start with the optimal choice and the steady state distribution generated by EconCast will converge to the optimal solution to (P4).

Next, to find the optimal Lagrange multipliers $\boldsymbol{\eta}^*$, consider the dual $\mathcal{D}(\boldsymbol{\eta}) := \mathcal{L}(\boldsymbol{\pi}^\eta, \boldsymbol{\eta})$ over $\boldsymbol{\eta} \succeq \mathbf{0}$ (here $\mathbf{0}$ is an N -dimensional zero vector and \succeq denotes component-wise inequality). Interestingly, it can be shown that the partial derivative of $\mathcal{D}(\boldsymbol{\eta})$ with respect to η_i is simply given by

$$\partial \mathcal{D} / \partial \eta_i = \rho_i - (\alpha_i L_i + \beta_i X_i), \quad (22)$$

which is the difference between the power budget ρ_i and the average power consumption of node i . Therefore, the dual can be minimized by using a gradient descent algorithm with inputs of step size $\delta_k > 0, \rho, \mathbf{L}$, and \mathbf{X} , which generates a state probability $\boldsymbol{\pi}(k)$ ($k = 1, 2, \dots$). This algorithm is described in Algorithm 1 along with the following equations

$$\eta_i(k) = [\eta_i(k-1) - \delta_k(\rho_i - \alpha_i(k)L_i - \beta_i(k)X_i)]^+, \quad (23)$$

$$\alpha_i(k) = \sum_{\mathbf{w} \in \mathcal{W}_i^L} \pi_{\mathbf{w}}^{\boldsymbol{\eta}(k)}, \quad \beta_i(k) = \sum_{\mathbf{w} \in \mathcal{W}_i^X} \pi_{\mathbf{w}}^{\boldsymbol{\eta}(k)}. \quad (24)$$

Hence, with the right choice of step size δ_k (e.g., $\delta_k = 1/k$), $\boldsymbol{\pi}(k)$ converges to the optimal solution to (P4).

Finally, to arrive at a distributed solution, instead of computing the quantities α_i and β_i directly according to (24) (which is centralized with high complexity), we can approximate the difference between the power budget and the average power consumption (22) by observing the dynamics of the energy storage level at each node. Specifically, each node i can update its Lagrange multiplier $\eta_i(k)$ based on the difference between its energy storage levels at the end and the start of an interval of length τ_k , divided by τ_k , as described by (17). Therefore η_i is updated according to a “noisy” gradient descent. However, it follows from stochastic approximation (with Markov modulated noise) that by choosing step sizes and interval lengths as given in Theorem 1, these noisy updates will converge to $\boldsymbol{\eta}^*$ as $k \rightarrow \infty$ (see e.g., Theorem 1 of [24]). As mentioned in Section 5.6, the choice of parameters σ, δ_k , and τ_k will affect the tradeoff between convergence time and the performance of EconCast.

7. NUMERICAL RESULTS

In this section, we consider various heterogeneous networks and numerically show that the throughput approaches the oracle throughput \mathcal{T}^* as σ decreases. In the special case of homogeneous networks, we explore the sensitivity of the

throughput to various power consumption levels, and compare to related work. Then, via simulations, we study the burstiness and latency of EconCast, and evaluate its performance in non-clique topologies.

Our general conclusions with respect to anyput performance are quite similar as to groupput. Therefore, we focus on the *groupput performance achieved by EconCast-C* throughout this section. For brevity, we omit the subscript for groupput and use the following notation: (i) \mathcal{T}^* is the oracle groupput obtained by solving (P1) or, equivalently, (P2), (ii) \mathcal{T}^σ is the achievable groupput of EconCast with a given value of σ obtained by solving (P4), and (iii) $\widetilde{\mathcal{T}}^\sigma$ is the groupput of EconCast obtained via simulations with a given value of σ .

7.1 Setup

We consider $\sigma \in \{0.1, 0.25, 0.5\}$. The nodes’ power budgets and consumption levels correspond to energy harvesting budgets and ultra-low-power transceivers in [21, 22, 45]. Unless stated otherwise, we use a power budget of $\rho = 10 \mu\text{W}$ and power consumption levels $L = X = 0.5 \text{ mW}$. This results in a ratio of 50 between the transceiver power consumption and the budget. Note that the performance of EconCast only depends on the *ratio* between the listen or transmit power and the power budget. For example, nodes with $\rho = 10 \mu\text{W}, L = X = 0.5 \text{ mW}$ behave exactly the same as nodes with $\rho = 1 \text{ mW}, L = X = 50 \text{ mW}$. Therefore, the oracle throughput applies and EconCast can operate in very general settings.

Recall that we assume that there are not simultaneous transmissions and collisions. We also assume that the packet length is 1 ms and that nodes have accurate estimate of the number of listeners, i.e., $\hat{c}(t) = c(t)$.

Our simulation results show that $\widetilde{\mathcal{T}}^\sigma$ perfectly matches \mathcal{T}^σ for $\sigma \in \{0.25, 0.5\}$. For $\sigma = 0.1$, $\widetilde{\mathcal{T}}^\sigma$ does not converge to \mathcal{T}^σ within reasonable time due to the bursty nature of EconCast, as will be described in Section 7.4. Therefore, we evaluate the throughput performance of EconCast by comparing \mathcal{T}^σ to \mathcal{T}^* with varying σ in both heterogeneous and homogeneous networks. Specifically, *homogeneous networks* consist of nodes with the same power budget and consumption levels, i.e., $\rho_i = \rho, L_i = L, X_i = X, \forall i \in \mathcal{N}$.

7.2 Heterogeneous Networks – Throughput

One strength of EconCast is its ability to deal with *heterogeneous* networks. Figure 2 shows the throughput achieved by EconCast normalized to the corresponding oracle throughput (i.e., $\mathcal{T}^\sigma / \mathcal{T}^*$) for heterogeneous networks with $N = 5$ and $\sigma \in \{0.1, 0.25, 0.5\}$. The results are obtained by solving (P1) for \mathcal{T}^* and (P4) for \mathcal{T}^σ with a given value of σ . Intuitively, higher values of $\mathcal{T}^\sigma / \mathcal{T}^*$ indicate better performance of EconCast.

Along the x -axis, the *network heterogeneity*, denoted by h , is varied from 10 to 250 at discrete points. The relationship between the network heterogeneity and the values of h is as follows: (i) for each node i , L_i and X_i are indepen-




Figure 2: Sensitivity of the achievable throughput normalized to the oracle throughput, $\mathcal{T}^\sigma/\mathcal{T}^*$, to the heterogeneity of the power budget, ρ , and consumption levels, L and X .

dently selected from a uniform distribution on the interval $[510 - h, 490 + h]$ (μW), (ii) for each node i , a variable h' is first sampled from the interval $[-\log \frac{h}{100}, \log h]$ uniformly at random, and then ρ_i is set to be $\exp(h')$. Therefore, the energy budget ρ_i varies from $100/h$ to h (μW). As a result, for any h , L_i and X_i have mean values of 0.5 mW. ρ_i has median of 10 μW but its mean increases as h increases. Note that a homogeneous network is represented by $h = 10$.

The y -axis indicates for each value of h , the mean and the 95% confidence interval of the ratios $\mathcal{T}^\sigma/\mathcal{T}^*$ averaged over 1000 heterogeneous network samples. Figure 2 shows that $\mathcal{T}^\sigma/\mathcal{T}^*$ approaches 1 as σ decreases, illustrating the results in Section 6. Furthermore, with increased heterogeneity of the network, $\mathcal{T}^\sigma/\mathcal{T}^*$ has little dependency on the heterogeneity but heavy dependency on σ . Since larger h values impose higher power budgets, the corresponding oracle throughput \mathcal{T}^* increases as well.

7.3 Homogeneous Networks – Comparison to Related Work

We now evaluate the performance of EconCast in *homogeneous* networks with respect to different power consumption levels, and compare to related work which assumes homogeneity across nodes. We consider three protocols: *Panda* [42], *Birthday* [43], and *Searchlight* [5], which operate under stricter assumptions than EconCast. In particular:

- The probabilistic protocols Panda and Birthday both require a homogeneous set of nodes and a priori knowledge of the number of nodes, N . The throughput of Panda and Birthday is computed as described in [42] and [43], respectively.
- The deterministic protocol Searchlight is designed for minimizing the worst case pairwise discovery latency, which does not directly address multi-party communication across a shared medium. However, the discovery latency is closely related to the throughput, since the inverse of the average latency is the throughput. Hence, maximizing throughput is equivalent to minimizing the average discovery latency. We derive an upper bound on




Figure 3: Throughput performance of various protocols (all normalized to the oracle throughput \mathcal{T}^*) with $N = 5$, $\rho = 10$ μW , and $L + X = 1$ mW, as a function of X/L .

the throughput of Searchlight by multiplying the pairwise throughput by $(N - 1)$. This is assuming that all other $(N - 1)$ nodes will be receiving when one node transmits. However, in practice the throughput is likely to be lower unless all the nodes are synchronized and coordinated.

Figure 3 presents the throughput achieved by various protocols normalized to the oracle throughput \mathcal{T}^* as a function of the ratio X/L , with $N = 5$, $\rho = 10$ μW , and $L + X = 1$ mW. The protocols considered are: EconCast with $\sigma \in \{0.1, 0.25, 0.5\}$, Panda, Birthday, and Searchlight.⁵ The horizontal dashed line at 1 represents the oracle throughput. Note that with $L = X = 0.5$ mW, the ratio $\mathcal{T}^\sigma/\mathcal{T}^*$ achieved by EconCast outperforms that of Panda by 6x and 17x with $\sigma = 0.5$ and $\sigma = 0.25$, respectively. The simulation results, which will be discussed later, also verify this throughput improvement.

Figure 3 shows that with $\sigma = 0.1$, \mathcal{T}^σ is very close to the oracle throughput \mathcal{T}^* unless $X \ll L$. As σ is increased, the ratio $\mathcal{T}^\sigma/\mathcal{T}^*$ drops as expected (see Section 6) but still significantly outperforms that of prior art for $X \approx L$. However, the performance of EconCast degrades with extreme values of X/L . This is because with small X/L values, nodes enter transmit state infrequently, since listen is expensive and they must pass the listen state to enter the transmit state. On the other hand, with large X/L values, nodes waste their energy to transmit even when there is no other nodes listening (e.g., $\hat{c}(t) = 0$). We believe that any distributed protocol will suffer from such performance degradation since, unlike Panda, Birthday, and Searchlight, nodes in a fully distributed setting *do not* have any information about the properties of other nodes in the network.

7.4 Burstiness and Latency

The results until now suggest allowing $\sigma \rightarrow 0$. While reducing σ improves throughput, it considerably increases the communication *burstiness*, as described in Section 5. In general, increased burstiness means that the long term throughput can be achieved with given power budgets but the vari-

⁵For Searchlight we compare its throughput upper bound to \mathcal{T}^* , as described in Section 7.3.




Figure 4: Average burst length obtained analytically (curves) and via simulation (markers) with $N \in \{5, 10\}$, $\sigma \in \{0.25, 0.5\}$, $\rho = 10 \mu\text{W}$, and $L = X = 500 \mu\text{W}$.




Figure 5: The CDF, mean, and 99th percentile latency of EconCast obtained via simulation with $N \in \{5, 10\}$, $\sigma \in \{0.25, 0.5\}$, $\rho = 10 \mu\text{W}$, and $L = X = 500 \mu\text{W}$, compared to the pairwise worst case latency of Searchlight.

ance is more significant during short term intervals. Recall from Section 7.1 that the packet length is 1 ms, we therefore measure the average burst length compared to this unit packet length.

Figure 4 shows the average burst length received by the nodes in homogeneous networks with $N \in \{5, 10\}$, $\rho = 10 \mu\text{W}$, $L = X = 500 \mu\text{W}$, and varying σ . Values are extracted using an analytical formula (curves) derived from the solution of (P4) and contrasted with simulations at specific values of σ (markers). Aside from showing that the simulation results and the analytical results are well matched, Figure 4 also demonstrates how reducing σ dramatically increases burstiness. For example, with $\sigma = 0.25$ and $N = 10$, a node has an average received burst length of 85 ms, and this value is increased to 4.5×10^5 ms with $\sigma = 0.1$. This explains why $\widetilde{\mathcal{T}}^\sigma$ cannot be obtained with $\sigma = 0.1$ (see Section 7.1) and we remark that reducing the communication burstiness is a subject of future work.

A second metric we consider is the communication *latency*. It is defined as the time interval between consecutive bursts received by a node from some other node where the interval includes *at least* one sleep period. We focus on this metric because nodes receiving longer bursts consume




Figure 6: The oracle throughput in non-clique topologies, \mathcal{T}_{nc}^* , and the throughput of EconCast obtained via simulations, $\widetilde{\mathcal{T}}^\sigma$, in grid topologies with varying N and σ .

more energy, and therefore, need to sleep for longer periods of time. It is possible for a node to receive multiple bursts in one listen period. Figure 5 presents the CDF of communication latency obtained via simulations, for $N \in \{5, 10\}$ and $\sigma \in \{0.25, 0.5\}$, and indicates both the average and the 99th-percentile latency values. It also shows the pairwise worst case latency of Searchlight computed from [5] under the same power consumption level and power budget.⁶

Figure 5 shows that (i) as σ decreases, the latency increases since nodes suffering from long bursts will sleep longer, and (ii) larger value of N results in lower latency, since it is more likely to receive when more nodes exist. As an example, with $\sigma = 0.5$ and $N = 5$, a node receives bursts from some other node on average every 5 seconds. In addition, for all parameters considered, the 99th-percentile latency is within 120 seconds, outperforming the Searchlight pairwise worst case latency bound of 125 seconds. Note that although EconCast has a non-zero probability of having any latency, in most cases (over 99%), its latency is below the worst case latency of Searchlight.

7.5 Evaluation in Non-clique Topologies

We now compute the oracle throughput for non-clique topologies (derived in Section 4.3) and evaluate the throughput of EconCast in such scenarios. Operating in a distributed manner, EconCast can be easily adapted to non-clique topologies: if a node hears two simultaneous transmissions from two nodes, it does not count any of the transmissions as throughput. Recall from Section 3 that nodes always check the channel utilization before waking up, simultaneous transmission can only happen if two nodes are not within communication range of each other.

We use grid topologies with varying number of nodes, N , in which a node can only have at most 4 neighbors. For example, $N = 25$ represents a 5×5 grid. For each value of N , we compute the oracle throughput (groupput). Figure 6 presents the oracle throughput, \mathcal{T}_{nc}^* , for grid topologies, and the throughput achieved by EconCast via simulations with

⁶This is computed with slot length of 50 ms and a beacon (packet) length of 1 ms as was done in [49].

with varying σ . Note that for all the grid topologies considered, the upper and lower bounds of \mathcal{T}_{nc}^* (see Section 4.3) are the same, providing the exact oracle throughput.

Figure 6 shows that EconCast achieves 14% – 22% of the maximum throughput \mathcal{T}_{nc}^* with $\sigma = 0.25$. Although increasing σ leads to lower throughput, it can be observed that as N increases, the throughput approaches 10% of \mathcal{T}_{nc}^* with $\sigma = 0.5$. Although we cannot obtain the throughput for $\sigma = 0.1$, achieving 10% – 20% of \mathcal{T}_{nc}^* is remarkable given the fact that EconCast works in a distributed manner in which each node does not have any information of the properties of other nodes.

8. EXPERIMENTAL EVALUATION

To experimentally evaluate the performance of EconCast-C,⁷ we implement it using the Texas Instruments eZ430-RF2500-SEH node [1].⁸ In this section, we first describe the energy measurements performed on the nodes running EconCast. Then, we describe the method by which nodes can estimate the number of listening nodes. Finally, we experimentally evaluate the performance of EconCast.

8.1 Experimental Setup

The TI eZ430-RF2500-SEH node is equipped with: (i) an ultra-low-power MSP430 microcontroller and a CC2500 wireless transceiver operating at 2.4 GHz at 250 Kbps, (ii) a solar energy harvester (SEH-01) that converts ambient light into electrical energy, and (iii) a 1 mF capacitor to power up the transceiver board. Despite its drawbacks which will be discussed below, it can be used for evaluation by extending the length of the shortest allowable data transmission.

We consider power budgets of $\rho \in \{1 \text{ mW}, 5 \text{ mW}\}$. From our measurements, a node spends $L = 67.08 \text{ mW}$ in the listen state and $X = 56.29 \text{ mW}$ in the transmit state.⁹ The power consumption levels are very similar from node to node. Recall from Section 7 that the performance of EconCast depends on the ratio between the power consumption levels and budget. Therefore, our experimental results will be similar to experiments when both the power consumption levels and budget are scaled down (e.g., a network of nodes with $\rho \in \{10 \mu\text{W}, 50 \mu\text{W}\}$, $L = 0.67 \text{ mW}$, and $X = 0.56 \text{ mW}$).

Each node is programmed with its ρ , L , and X as the input of EconCast-C. The nodes' main drawbacks include (i) inaccurate readings of the energy storage level (i.e., the voltage of the on-board capacitor) which are sensitive to the environment, and (ii) the fact that the 1 mF capacitor cannot support multiple packet transmissions. Due to these drawbacks, we implement (via software) a virtual battery at each node. The virtual battery emulates the node's energy storage level based on its sleep, listen, and transmit activities,

⁷See Section 5.4 the reasons for only implementing EconCast-C.

⁸A demonstration of the testbed is presented in [11].

⁹This corresponds to a -16 dBm transmission power, at which nodes within the same room typically have little or no packet loss.

and is used for updating the Lagrange multiplier according to (17). We show in the following section that in practice, a node running EconCast-C using this virtual battery is indeed consuming power at a rate close to its power budget.

8.2 Energy Consumption Measurements

To accurately measure the power consumption of the nodes, we disable the on-board solar cell, and attach a large pre-charged capacitor ($C_{cap} = 5 \text{ F}$) that stores energy in advance (similar power consumption measurements were used in [55]). The energy consumed is computed by

$$E_{consumed} = 0.5C_{cap} \cdot (V_{t_0}^2 - V_{t_1}^2), \quad (25)$$

where V_{t_0} and V_{t_1} are the measured power voltage values of the capacitor at t_0 and t_1 . The empirical average power consumption, \bar{P} (mW), is then computed by

$$\bar{P} = E_{consumed} / (t_1 - t_0). \quad (26)$$

Note that even with such a big capacitor, a node with a power budget of 1 mW (5 mW) has a lifetime of only 135 (27) minutes with $V_{t_0} = 3.6 \text{ V}$ and $V_{t_1} = 3.0 \text{ V}$, which represent its stable working voltage range.

To measure the power consumption of the nodes, we charge the capacitor to $V_{t_0} = 3.6 \text{ V}$ and log the readings of V_{t_1} after 30 minutes using a multimeter. The empirical average power consumption is computed from (25) and (26) for $\sigma \in \{0.25, 0.5\}$ and is averaged using 60 runs. Because L and X do not account for some additional energy usage,¹⁰ the *actual power consumption*, \bar{P} , is in fact a small fraction higher than the *target power budget*, ρ . Irrespective of σ , the measurement results show that \bar{P} exceeds ρ by 11% for $\rho = 1 \text{ mW}$, and by 4% for $\rho = 5 \text{ mW}$.

Observing the empirical power consumption of the nodes, we compute the achievable throughput by solving (P4) using both the actual power consumption, \bar{P} , and the target power budget, ρ , denoted by $\bar{\mathcal{T}}^\sigma$ and \mathcal{T}^σ , respectively. In Section 8.4, we compare the experimental throughput to both $\bar{\mathcal{T}}^\sigma$ and \mathcal{T}^σ . Having verified the power consumption of the nodes, we replace the capacitor with AAA batteries,¹¹ allowing the experiments to run for longer times.

8.3 Practical Pinging

To enable practical pinging in EconCast-C, a short, fixed-length *pinging interval* is introduced after each packet transmission. During this interval, the transmitter listens for pings and recipients of the previous packet send a short ping at a random time uniformly distributed within the interval. The transmitter then estimates the number of listeners, $\hat{c}(t)$, by counting the pings it receives, and adjusts the transition rate, $\lambda_{xi}^C(t)$, according to (18e).

Ideally, each ping should be much shorter than both the pinging interval and the packet length in order to reduce the collisions between pings, as well as for the transmitter to

¹⁰The additional energy usage includes the energy consumed in powering up the regulator circuitry, etc.

¹¹The constant voltage of AAA batteries limits the ability to measure the power consumption of the nodes.




Figure 7: Points marked as “Ideal” (“Relaxed”) represent ratio of experimental throughput normalized to the achievable throughput obtained by using the target power budget (actual power consumption) and points marked as “Battery Variance” present the average, minimum, and maximum ratios of power consumption normalized to target power budget, with $N \in \{5, 10\}$, $\rho \in \{1 \text{ mW}, 5 \text{ mW}\}$, and $\sigma \in \{0.25, 0.5\}$.

successfully receive it. Therefore, we use pings of length 0.4 ms, which is the shortest packet that can be sent by a node. Based on this, we empirically set the pinging interval to 8 ms and each data packet to 40 ms.

8.4 Performance Evaluation

We consider homogeneous networks of size $N \in \{5, 10\}$, power budgets $\rho \in \{1 \text{ mW}, 5 \text{ mW}\}$, $\sigma \in \{0.25, 0.5\}$, and nodes that are all located in proximity. One additional listening node (a 6th or 11th node) is also present but only as an observer and is connected to a PC via a USB port. Each data packet contains the node ID and information about the number of packets it has received from each other node. The observer node reports all received packets to the PC for storage and post processing. Each experiment is conducted for up to 24 hours. The experimental throughput is computed by dividing the duration of successful transmissions by the experiment duration.

Throughput evaluation: Figure 7 presents the ratio of the experimentally obtained throughput, $\tilde{\mathcal{T}}^\sigma$, normalized to the achievable throughput \mathcal{T}^σ and $\overline{\mathcal{T}}^\sigma$ (see Section 8.2). Separate charts represent the results for differing power budget, ρ , and number of nodes, N . Points marked “Ideal” show the experimental throughput normalized to the achievable throughput computed by solving (P4) with the target power budget ρ (i.e., $\tilde{\mathcal{T}}^\sigma/\mathcal{T}^\sigma$). Points marked “Relaxed” show the experimental throughput normalized to the achievable throughput computed by solving (P4) with the actual power consumption \bar{P} (i.e., $\tilde{\mathcal{T}}^\sigma/\overline{\mathcal{T}}^\sigma$). As expected, $\overline{\mathcal{T}}^\sigma$ is higher than \mathcal{T}^σ , resulting in a lower throughput ratio.

Figure 7 shows that despite the practical limitations (e.g., packet collisions and inaccurate clocks) faced when running EconCast-C on real hardware, the ratio $\tilde{\mathcal{T}}^\sigma/\overline{\mathcal{T}}^\sigma$ is between 57% – 77% ($\tilde{\mathcal{T}}^\sigma/\mathcal{T}^\sigma$ is between 67% – 81%) for all settings considered. Moreover, Table 3 shows the improvement of EconCast-C over the throughput of Panda computed according to [42], denoted by $\mathcal{T}_{\text{Panda}}$, under the same power consumption levels and budget, with $\sigma = 0.25$. It can be seen that with power budget of $\rho = 1 \text{ mW}$, the experimental throughput of EconCast-C outperforms the analytically computed throughput of Panda by 8x – 11x.

$(N, \rho(\text{mW}))$	(1, 5)	(1, 10)	(5, 5)	(5, 10)
$\tilde{\mathcal{T}}^\sigma/\mathcal{T}^\sigma$ (%)	66.78	77.96	74.84	80.53
$\mathcal{T}_{\text{Panda}}/\mathcal{T}^\sigma$ (%)	6.24	9.64	19.35	35.63
$\tilde{\mathcal{T}}^\sigma/\mathcal{T}_{\text{Panda}}$	10.76	8.09	3.87	2.26

Table 3: Experimental throughput of EconCast-C compared to computed throughput of Panda (all normalized to the achievable throughput \mathcal{T}^σ), with $\sigma = 0.25$ and varying (N, ρ) .

We remark that getting a higher experimental throughput ratio is limited by the following reasons. First, there is an 8 ms pinging interval (see Section 8.3) after each packet transmission which effectively reduces the number of bits delivered. Second, collisions of pings or failed decodings of pings result in inaccurate estimates of the number of listeners. Third, the low-power clock used by a node during its sleep state drifts and additionally can be affected by its environment.

Power consumption: Recall that in the power consumption measurements described in Section 8.2, we show that the power consumption of the virtual battery is valid for evaluating the actual power consumption of the node. Table 3 also presents the mean, minimum, and maximum power consumption of the virtual battery normalized to the target power budget ρ . Specifically, a value of 1 means that a node consumes power on average at the rate of its power budget throughout the experiment, and a higher value means that a node consumes power at the rate which is higher than its power budget.

The results show that nodes running EconCast-C consume power at rates which are within 7% and 3% of the target power budget with $\sigma = 0.25$ and $\sigma = 0.5$, respectively. This is because smaller value of σ increases the communication burstiness (see Section 7.4), resulting in larger variance of the nodes’ virtual battery levels.

Collection of Pings: An important input to EconCast-C is the estimates of number of active listeners, $\hat{c}(t)$, based on which the transmitter decides the probability to continuously transmit. Larger values of $\hat{c}(t)$ lead to longer average burst length and can potentially significantly increase the throughput. For example, receiving 1 ping, the transmit-

# of Listeners	0	1	2	3	4
$\rho = 1 \text{ mW}(\%)$	89.03	9.69	1.28	0.00	0.00
$\rho = 5 \text{ mW}(\%)$	59.21	31.22	8.22	1.24	0.11

Table 4: Distribution of number of pings (active listeners) received after each packet transmission with $N = 5$, $\sigma = 0.25$, and varying ρ .

ter continuously transmits a packet with probability 0.8647 with $\sigma = 0.5$. This probability increases to 0.9817 with $\sigma = 0.25$, which substantially increases the burstiness. Also, with lower power budget, a successful transmission happens more rarely and it becomes harder to collect pings.

Table 4 presents the distribution of number of pings (equivalently, number of active listeners) received by the transmitter after each packet transmission, during experiments of $N = 5$, $\sigma = 0.25$, and $\rho \in \{1 \text{ mW}, 5 \text{ mW}\}$. It can be shown that with a higher power budget, the nodes are more active and the transmitter has higher probability to receive more pings. On the other hand, with lower power budget, the transmitter almost never receives more than 3 pings in a 5 nodes experiment, resulting in lower throughput as illustrated in Figure 7.

9. CONCLUSION

In this paper, we considered the problem of maximizing the broadcast groupput and anyput among a set of energy-constrained nodes with heterogeneous power budgets and listen and transmit power consumption levels. We provided methods to obtain oracle groupput and oracle anyput for a given set of heterogeneous nodes.

We developed the EconCast-C and EconCast-NC distributed protocols that control the nodes' transitions among sleep, listen, and transmit states. We analytically showed that heterogeneous nodes using the protocols (without any a priori knowledge regarding the number of nodes, power consumption levels, and budgets) can achieve the oracle groupput and anyput in a limiting sense (when $\sigma \rightarrow 0$).

Through simulations we evaluated EconCast and compared it to the state of the art. We also considered the trade-offs in its design as a function of σ , where low values increase both burstiness and throughput, while high values reduce both of them. Finally, we experimentally evaluated EconCast using commercial-off-the-shelf nodes, thereby demonstrating its practicality.

There are several open future research directions. In particular, future research will focus on extending the analysis to non-clique topologies. Moreover, evaluation with custom designed ultra-low-power nodes (e.g., [41]), that have improved energy awareness compared to the TI eZ430-RF2500-SEH nodes, would enable to better assess the trade-offs related to the protocol design. Finally, considering unique application characteristics and their relation to groupput and anyput is an open problem.

10. ACKNOWLEDGMENTS

This research was supported in part by ARO grant 9W911NF-16-1-0259, NSF grant ECCS-1547406, and the

People Programme (Marie Curie Actions) of the European Union's Seventh Framework Programme (FP7/2007-2013) under REA grant agreement n°[PIIF-GA-2013-629740].11.

11. REFERENCES

- [1] Texas Instruments eZ430-RF2500-SEH solar energy harvesting development tool user's guide, 2013.
- [2] Texas Instruments CC2640 SimpleLink Bluetooth smart wireless MCU, 2015.
- [3] L. Atzori, A. Iera, and G. Morabito. The internet of things: A survey. *Computer Networks*, 54(15):2787 – 2805, 2010.
- [4] W. Bajwa, J. Haupt, A. Sayeed, and R. Nowak. Compressive wireless sensing. In *Proc. ACM/IEEE IPSN'06*, 2006.
- [5] M. Bakht, M. Trower, and R. H. Kravets. Searchlight: Won't you be my neighbor? In *Proc. ACM MobiCom'12*, 2012.
- [6] S. A. Borbash, A. Ephremides, and M. J. McGlynn. An asynchronous neighbor discovery algorithm for wireless sensor networks. *Ad Hoc Netw.*, 5(7):998 – 1016, 2007.
- [7] M. Buettner, B. Greenstein, and D. Wetherall. Dewdrop: an energy-aware runtime for computational RFID. In *Proc. USENIX NSDI'11*, 2011.
- [8] L. Chen, R. Fan, K. Bian, M. Gerla, T. Wang, and X. Li. On heterogeneous neighbor discovery in wireless sensor networks. *arXiv preprint arXiv:1411.5415*, 2014.
- [9] L. Chen, Y. Gu, S. Guo, T. He, Y. Shu, F. Zhang, and J. Chen. Group-based discovery in low-duty-cycle mobile sensor networks. In *Proc. IEEE SECON'12*, 2012.
- [10] S. Chen, P. Sinha, N. B. Shroff, and C. Joo. A simple asymptotically optimal joint energy allocation and routing scheme in rechargeable sensor networks. *IEEE/ACM Trans. Netw.*, 22(4):1325–1336, 2014.
- [11] T. Chen, G. Chen, S. Jain, R. Margolies, G. Grebla, D. Rubenstein, and G. Zussman. Demo abstract: Power-aware neighbor discovery for energy harvesting things. In *Proc. ACM SenSys'16*, 2016.
- [12] T. Chen, J. Ghaderi, D. Rubenstein, and G. Zussman. Maximizing broadcast throughput under ultra-low-power constraints. *arXiv preprint arXiv:1610.04203*, 2016.
- [13] M. Chiang, S. H. Low, J. C. Doyle, et al. Layering as optimization decomposition: A mathematical theory of network architectures. *Proc. IEEE*, 95(1):255–312, 2007.
- [14] S. Cui, A. J. Goldsmith, and A. Bahai. Energy-efficiency of MIMO and cooperative MIMO techniques in sensor networks. *IEEE J. Sel. Areas Commun.*, 22(6):1089–1098, 2004.
- [15] A. G. Dimakis, V. Prabhakaran, and K. Ramchandran. Ubiquitous access to distributed data in large-scale sensor networks through decentralized erasure codes. In *Proc. ACM/IEEE IPSN'05*, 2005.

- [16] P. Dutta and D. Culler. Practical asynchronous neighbor discovery and rendezvous for mobile sensing applications. In *Proc. ACM SenSys'08*, 2008.
- [17] A. El-Hoiydi and J.-D. Decotignie. Wisemac: an ultra low power MAC protocol for the downlink of infrastructure wireless sensor networks. In *Proc. IEEE ISCC'04*, 2004.
- [18] J. Ghaderi, S. Borst, and P. Whiting. Queue-based random-access algorithms: Fluid limits and stability issues. *Stochastic Systems*, 4(1):81–156, 2014.
- [19] J. Ghaderi and R. Srikant. On the design of efficient CSMA algorithms for wireless networks. In *Proc. IEEE CDC'10*, 2010.
- [20] M. Gorlatova, P. Kinget, I. Kymissis, D. Rubenstein, X. Wang, and G. Zussman. Challenge: ultra-low-power energy-harvesting active networked tags (EnHANTs). In *Proc. ACM MobiCom'09*, 2009.
- [21] M. Gorlatova, J. Sarik, G. Grebla, M. Cong, I. Kymissis, and G. Zussman. Movers and shakers: Kinetic energy harvesting for the internet of things. *IEEE J. Sel. Areas Commun.*, 33(8):1624–1639, 2015.
- [22] M. Gorlatova, A. Wallwater, and G. Zussman. Networking low-power energy harvesting devices: Measurements and algorithms. *IEEE Trans. Mobile Comput.*, 12(9):1853–1865, Sept. 2013.
- [23] J. Haupt, W. U. Bajwa, M. Rabbat, and R. Nowak. Compressed sensing for networked data. *IEEE Sigal Process Mag.*, 25(2):92–101, 2008.
- [24] L. Jiang and J. Walrand. Convergence and stability of a distributed CSMA algorithm for maximal network throughput. In *Proc. CDC'09*, 2009.
- [25] L. Jiang and J. Walrand. A distributed CSMA algorithm for throughput and utility maximization in wireless networks. *IEEE/ACM Trans. Netw.*, 18:960–972, 2010.
- [26] A. Kandhalu, K. Lakshmanan, and R. R. Rajkumar. U-connect: A low-latency energy-efficient asynchronous neighbor discovery protocol. In *Proc. ACM/IEEE IPSN'10*, 2010.
- [27] S. Katti, S. Gollakota, and D. Katabi. Embracing wireless interference: Analog network coding. In *Proc. ACM SIGCOMM'07*, 2007.
- [28] L. Keller, E. Atsan, K. Argyraki, and C. Fragouli. Sensecode: Network coding for reliable sensor networks. *ACM Trans. Sens. Netw.*, 2013.
- [29] L. Keller, M. J. Siavoshani, C. Fragouli, K. Argyraki, and S. Diggavi. Identity aware sensor networks. In *Proc. IEEE INFOCOM'09*, 2009.
- [30] F. P. Kelly, A. K. Maulloo, and D. K. Tan. Rate control for communication networks: shadow prices, proportional fairness and stability. *J. Operational Research Society*, 49(3):237–252, 1998.
- [31] M. Khouzani, S. Eshghi, S. Sarkar, N. B. Shroff, and S. S. Venkatesh. Optimal energy-aware epidemic routing in DTNs. In *Proc. ACM MobiHoc'12*, 2012.
- [32] J.-S. Lee, Y.-W. Su, and C.-C. Shen. A comparative study of wireless protocols: Bluetooth, UWB, ZigBee, and Wi-Fi. In *Proc. IEEE IECON'07*, 2007.
- [33] J. Liu, Y. Yi, and A. Proutiere. Towards utility-optimal random access without message passing. *Wireless Commun. Mobile Comput.*, 10:115–128, 2010.
- [34] V. Liu, A. Parks, V. Talla, S. Gollakota, D. Wetherall, and J. R. Smith. Ambient backscatter: wireless communication out of thin air. In *Proc. ACM SIGCOMM'13*, 2013.
- [35] S. H. Low and D. E. Lapsley. Optimization flow control - I: basic algorithm and convergence. *IEEE/ACM Trans. Netw.*, 7(6):861–874, 1999.
- [36] S. H. Low and R. Srikant. A mathematical framework for designing a low-loss, low-delay internet. *Networks and Spatial Economics*, 4(1):75–101, 2004.
- [37] C. Luo, F. Wu, J. Sun, and C. W. Chen. Compressive data gathering for large-scale wireless sensor networks. In *Proc. ACM MobiCom'09*, 2009.
- [38] J. Luo, L. Xiang, and C. Rosenberg. Does compressed sensing improve the throughput of wireless sensor networks? In *Proc. IEEE ICC'10*, 2010.
- [39] Z. Mao, C. E. Koksal, and N. B. Shroff. Cross-layer resource allocation in energy harvesting sensor networks. In *Rechargeable Sensor Networks: Technology, Theory, and Application: Introducing Energy Harvesting to Sensor Networks*. 2014.
- [40] P. Marbach and A. Eryilmaz. A backlog-based CSMA mechanism to achieve fairness and throughput-optimality in multihop wireless networks. In *Proc. Allerton Conf.*, 2008.
- [41] R. Margolies, M. Gorlatova, J. Sarik, G. Stanje, J. Zhu, P. Miller, M. Szczodrak, B. Vignaham, L. Carloni, P. Kinget, et al. Energy-harvesting active networked tags (EnHANTs): Prototyping and experimentation. *ACM Trans. Sens. Netw.*, 11(4):62, 2015.
- [42] R. Margolies, G. Grebla, T. Chen, D. Rubenstein, and G. Zussman. Panda: Neighbor discovery on a power harvesting budget. In *Proc. IEEE INFOCOM'16*, 2016.
- [43] M. J. McGlynn and S. A. Borbash. Birthday protocols for low energy deployment and flexible neighbor discovery in ad hoc wireless networks. In *Proc. ACM MobiHoc'01*, 2001.
- [44] P. V. Nikitin, S. Ramamurthy, R. Martinez, and K. Rao. Passive tag-to-tag communication. In *Proc. IEEE RFID'12*, 2012.
- [45] K. Philips. Ultra low power short range radios: Covering the last mile of the IoT. In *Proc. IEEE ESSCIRC'14*, 2014.
- [46] A. Purohit, B. Priyantha, and J. Liu. WiFlock: Collaborative group discovery and maintenance in mobile sensor networks. In *Proc. ACM/IEEE IPSN'11*, 2011.

- [47] M. Rabbat, J. Haupt, A. Singh, and R. Nowak. Decentralized compression and predistribution via randomized gossiping. In *Proc. ACM/IEEE IPSN'06*, 2006.
- [48] T. Srisooksai, K. Keamarungsi, P. Lamsrichan, and K. Araki. Practical data compression in wireless sensor networks: A survey. *J. Netw. Comput. Appl.*, 35(1):37–59, 2012.
- [49] W. Sun, Z. Yang, K. Wang, and Y. Liu. Hello: A generic flexible protocol for neighbor discovery. In *Proc. IEEE INFOCOM'14*, 2014.
- [50] W. Sun, Z. Yang, X. Zhang, and Y. Liu. Energy-efficient neighbor discovery in mobile ad hoc and wireless sensor networks: A survey. *IEEE Commun. Surveys Tuts.*, 16(3):1448–1459, 2014.
- [51] S. Ulukus, A. Yener, E. Erkip, O. Simeone, M. Zorzi, P. Grover, and K. Huang. Energy harvesting wireless communications: A review of recent advances. *IEEE J. Sel. Areas Commun.*, 33(3):360–381, Mar. 2015.
- [52] S. Vasudevan, M. Adler, D. Goeckel, and D. Towsley. Efficient algorithms for neighbor discovery in wireless networks. *IEEE/ACM Trans. Netw.*, 21(1):69–83, 2013.
- [53] S. Vasudevan, J. Kurose, and D. Towsley. On neighbor discovery in wireless networks with directional antennas. In *Proc. IEEE INFOCOM'05*, 2005.
- [54] S. Vasudevan, D. Towsley, D. Goeckel, and R. Khalili. Neighbor discovery in wireless networks and the coupon collector's problem. In *Proc. ACM MobiCom'09*, Sept. 2009.
- [55] J. Wang, H. Hassanieh, D. Katabi, and P. Indyk. Efficient and reliable low-power backscatter networks. In *Proc. ACM SIGCOMM'12*. ACM, 2012.
- [56] J. Wang and D. Katabi. Dude, where's my card?: RFID positioning that works with multipath and non-line of sight. In *Proc. ACM SIGCOMM'13*, 2013.
- [57] Y. Wu, P. Chou, S.-Y. Kung, et al. Minimum-energy multicast in mobile ad hoc networks using network coding. *IEEE Trans. Commun.*, 53(11):1906–1918, 2005.
- [58] K. Xu, O. Dousse, and P. Thiran. Self-synchronizing properties of CSMA wireless multi-hop networks. In *Proc. ACM SIGMETRICS'10*, 2010.
- [59] W. Ye, J. Heidemann, and D. Estrin. Medium access control with coordinated adaptive sleeping for wireless sensor networks. *IEEE/ACM Trans. Netw.*, 12(3):493–506, 2004.
- [60] Y. Zeng, K. A. Mills, S. Gokhale, N. Mittal, S. Venkatesan, and R. Chandrasekaran. Robust neighbor discovery in multi-hop multi-channel heterogeneous wireless networks. *arXiv preprint arXiv:1505.00267*, 2015.