

Maximizing Broadcast Throughput Under Ultra-Low-Power Constraints

Tingjun Chen, Javad Ghaderi, Dan Rubenstein, and Gil Zussman

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. **EconCast** can operate in groupput or anyput modes to respectively maximize two alternative throughput measures. We show that **EconCast** approaches the oracle throughput. The performance is also evaluated numerically and via extensive simulations and it is shown that **EconCast** outperforms prior art by 6x – 17x under realistic assumptions. Moreover, we evaluate **EconCast**’s latency performance and consider design tradeoffs when operating in groupput and anyput modes. 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.

Index Terms—Internet-of-Things, energy harvesting, ultra-low-power, wireless communication

I. INTRODUCTION

Object tracking and monitoring applications are gaining popularity within the realm of Internet-of-Things [2]. 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 [3], [4]. Relying on node-to-node communications will require less infrastructure than traditional (RFID/reader-based) implementations. Therefore, as discussed in [3]–[6], it is envisioned that such ultra-low-power nodes will facilitate tracking applications

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. A partial and preliminary version of this paper appeared in ACM CoNEXT’16, Dec. 2016 [1], and some results were presented in IEEE WCNC’17 Workshop on Energy Harvesting and Remotely Powered Wireless Communications for the IoT (invited), Mar. 2017.

T. Chen, J. Ghaderi, and G. Zussman are with the Department of Electrical Engineering, Columbia University, New York, NY, USA (email: {tingjun, jghaderi, gil}@ee.columbia.edu).

D. Rubenstein is with the Department of Computer Science, Columbia University, New York, NY 10027, USA (e-mail: danr@cs.columbia.edu).

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 0.01–0.1 mW [7], [8] (for more details see the review in [9] 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 (BLE) [10] and ZigBee/802.15.4 [11] (usually at the order of 1–10 mW). On the other hand, BLE 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 [7]. 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 [12]–[14]. 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 [15], [16], or decentralized coding [17]. 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 groupput and anyput (i.e., maximum throughput achieved by an oracle) and provide methods to efficiently compute their values. Then, we use Lagrangian methods and a Q-CSMA (Queue-based Carrier Sense Multiple Access) approach to design EconCast. EconCast can operate in groupput or anyput modes to respectively maximize the two alternative throughput measures. Nodes running EconCast 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 the throughput performance of EconCast numerically and via extensive simulations under a wide range of power budgets and listen/transmit consumption levels, and for various heterogeneous and homogeneous networks. Specifically, numerical results show that EconCast outperforms prior art (Panda [14], Birthday [18], and Searchlight [19]) by a factor of 6x – 17x under realistic assumptions. In addition, we consider the performance of EconCast in terms of burstiness and latency. We also consider the design tradeoffs of EconCast when operating in groupput and anyput modes.

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 experimental throughput to analytical throughput of Panda [14] (where the analytical throughput is usually better than the experimental performance) and show that, for example, EconCast outperforms Panda by 8x – 11x.

We note that EconCast is not designed based on specific assumption, regarding the topology and that nodes do not need to know the properties of other nodes. Yet, in this paper, we mainly focus on a clique topology (i.e., all nodes are within the communication range of each other), since it lends itself to analysis. We briefly extend the analytical results to non-clique topologies and also evaluate the performance of EconCast in 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, and (iii) extensive performance evaluation of the protocol.

The rest of the paper is organized as follows. We discuss related work in Section II and formulate the problem in Section III. In Section IV, we present methods to efficiently compute the oracle throughput. We present EconCast in Section V and the proof of the main theoretical result in Section VI. In Section VII, we evaluate EconCast numerically and via simulations. We then discuss the experimental implementation and evaluation of EconCast in Section VIII. We conclude in Section IX.

II. RELATED WORK

There is vast amount of related literature on sensor networking and neighbor discovery that tries to limit energy consumption. Most of the protocols do not explicitly account for different listen and transmit power consumption levels of the nodes, or do not account for different power budgets [12], [13], [18]–[25]. 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 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, these protocols often require some explicit coordination (e.g., slotting [12], [18] or exchange of parameters [14], [19]), 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., [26]), and queue-based CSMA literature (e.g., [27]–[31]). 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., [32]) 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

TABLE I: Nomenclature

\mathcal{N}, N	Set of nodes, number of nodes
L_i, \mathbf{L}	Node i 's listen power consumption (Watt), $\mathbf{L} = [L_i]$
X_i, \mathbf{X}	Node i 's transmit power consumption (Watt), $\mathbf{X} = [X_i]$
ρ_i, ρ	Node i 's power budget (Watt), $\rho = [\rho_i]$
b_i	Energy storage level of node i (Joule)
\mathbf{w}, \mathcal{W}	Network state, the set of collision-free states
α_i, α	Fraction of time node i listens, $\alpha = [\alpha_i]$
β_i, β	Fraction of time node i transmits, $\beta = [\beta_i]$
$\gamma, \hat{\gamma}$	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
$\pi_{\mathbf{w}}, \pi$	Fraction of time the network is in $\mathbf{w} \in \mathcal{W}$, $\pi = [\pi_{\mathbf{w}}]$
$\mathcal{T}_{\mathbf{w}}$	Throughput of state $\mathbf{w} \in \mathcal{W}$
$\mathcal{T}, \mathcal{T}^*$	Throughput and oracle throughput
$\mathcal{T}_g, \mathcal{T}_g^*$	Groupput and oracle groupput
$\mathcal{T}_a, \mathcal{T}_a^*$	Anyput and oracle anyput
η_i, η	Lagrange multiplier of node i , $\eta = [\eta_i]$

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 technology such as backscatter [6].

III. 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 I summarizes the notations.

A. 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 levels are 0, L_i (W), and X_i (W).² These values 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 [7], [8] 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 $j \neq i$.

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

²The actual power consumption in *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 level.

B. 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 V how this property will help us develop a distributed protocol and in Section VIII, we provide practical means by which such estimates can be obtained.

C. 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 V) 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 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

³We also investigate non-clique networks in Section IV-C.

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.

D. 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

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

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

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

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

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 IV, 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 V, we focus on the design of EconCast that makes these decisions based on (P1).

IV. 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 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 where *at most* one transmitter can be present at any time imposes (11), which bounds the sum of the transmit fractions by 1.

A. 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.,

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

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

$$(P2) \quad \mathcal{T}_g^* := \max_{\boldsymbol{\alpha}, \boldsymbol{\beta}} \sum_{i \in \mathcal{N}} \alpha_i \quad (13)$$

$$\text{subject to } (9) - (12).$$

(P2) is an LP consisting of $2N$ variables and $(3N + 1)$ constraints (i.e., solving for $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ given inputs of N , $\boldsymbol{\rho}$, \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.

Lemma 1: The (rational-valued) solution $(\boldsymbol{\alpha}^*, \boldsymbol{\beta}^*)$ 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.

Proof: The proof can be found in Appendix A. ■

In *homogeneous* networks (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 solution to (P2) is given by⁴

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

⁴We show that in the optimal solution, the equalities hold for equations (9) and (12). The details are Appendix B.

B. Oracle Anput in a Clique

The oracle anput is obtained based on the observation that a transmission only occurs when there is *at least* one listener. Define variables $\chi_{i,j}$ as the fraction of time node j receives a transmission from node i , for the following 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 anput*, denoted by \mathcal{T}_a^* , can be obtained by solving the following maximization problem

$$(P3) \quad \mathcal{T}_a^* := \max_{\alpha, \beta} \sum_{i \in \mathcal{N}} \beta_i \quad (16)$$

subject to (9) – (11), (14), and (15).

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 anput by ensuring that every transmission is received by *at least* one node.

In homogeneous networks, the closed-form solution to (P3) is given by

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

C. 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

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

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 VII-E, we compute \mathcal{T}_{nc}^* and evaluate the performance of EconCast in non-clique topologies.

V. 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

TABLE II: A simple example in a heterogeneous network.

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

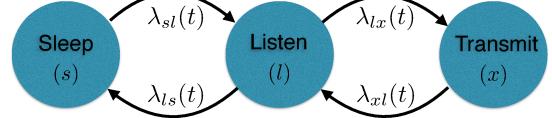


Fig. 1: The node's states and transition rates.

single node i , in parts of this section, we drop the subscript i of previously defined variables for notational compactness.

A. 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 power consumption $L_i = X_i = 1$ mW ($i = 1, 2, 3, 4$), but different power budgets ρ_i , as indicated in Table II. Table II 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., $\frac{100 \cdot \beta_i^*}{\alpha_i^* + \beta_i^*}\%$).

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).*

B. Protocol Description

To clearly present EconCast, we start from a theoretical framework and slowly build on it to address practicalities. As mentioned in Section III, a node can be in one of three states: sleep (s), listen (l), and transmit (x). As depicted in Fig. 1, it must pass through listen state to transition between sleep and transmit states. The time duration a node spends in 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.

To maximize the groupput or anput, EconCast can operate in *groupput mode* or *anput mode*, respectively. The throughput as a function of π_w (see (5)) 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, including

- Its power consumption levels in listen and transmit states, 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 (in groupput mode), $c(t)$, or an indicator of whether there are any active listeners (in anyput mode), $\gamma(t)$. 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.

Note that in EconCast, unlike in previous work such as Panda [14], each node *does not* need to know the number of nodes in the network, N , and the power budgets and 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 [4]), 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 V-E.

C. Estimating Active Listeners: Pings

We now discuss the estimation of $\hat{c}(t)$ or $\hat{\gamma}(t)$. Recall from Section III 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.

D. 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 V-E). In Section VIII-C, 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 V-E).

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 VIII focus on EconCast-C.

E. 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

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

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 senses the ongoing transmission).

The transition rates are described as follows (the superscripts C and N denote EconCast-C and EconCast-NC), in which the two throughput modes only differ in $\lambda_{xl}(t)$ (for the capture version) or in $\lambda_{lx}(t)$ (for the non-capture version). 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 and the proof is in Section VI.

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).

F. 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 decreased, the throughput increases, as we will describe in Section VI. However, the *burstiness* also increases with respect to decreased σ . The burstiness is a characteristic of communication involving multiple packets that are successfully received in bursts. In Section VII, we describe how the burstiness can be analyzed and measured.

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 .

VI. PROOF OF THEOREM 1

The proof of Theorem 1 is based on a Markov Chain Monte Carlo (MCMC) approach [28], [30] from statistical physics and consists of three parts: (i) we compute the steady state distribution of the network Markov chain under EconCast with fixed Lagrange multiplier vector $\eta = [\eta_i]$, (ii) we present an alternative concave optimization problem whose optimal value approaches that of (P1) as $\sigma \rightarrow 0$ and show that the steady state distribution of EconCast is indeed the optimal solution to this alternative optimization problem when the Lagrange multipliers are chosen optimally, and (iii) we show that under EconCast, nodes update their Lagrange multipliers locally according to a “noisy” gradient descent which converge to the optimal Lagrange multipliers with proper choices of step sizes and interval lengths as given in Theorem 1.

Part (i): Steady State Distribution

The following lemma describes the network state distribution generated by EconCast when η freezes.

Lemma 2: With fixed η , the network Markov chain, resulted from overall interactions among the nodes according to the transition rates (18), has the steady state distribution

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

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

Proof: The proof can be found in Appendix C. ■

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 $\pi(k)$ from (19) using $\eta = \eta(k)$
- 3: **for** $i = 1, 2, \dots, N$ **do**
- 4: Update $\eta_i(k)$, $\alpha_i(k)$, and $\beta_i(k)$ according to (23), (24)

Part (ii): An Alternative Optimization

We then present an optimization problem (P4) as follows

$$(P4) \max_{\pi} \sum_{\mathbf{w} \in \mathcal{W}} \pi_{\mathbf{w}} \mathcal{T}_{\mathbf{w}} - \sigma \sum_{\mathbf{w} \in \mathcal{W}} \pi_{\mathbf{w}} \log \pi_{\mathbf{w}} \quad (20)$$

subject to (6), (7), and (8),

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 $\mathcal{L}(\pi, \eta)$ formulated by moving the power constraint (6) into the objective (20) with a Lagrange multiplier $\eta_i \geq 0$ for each node i , i.e.,

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

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

In order to find η^* , consider the dual function $\mathcal{D}(\eta) = \mathcal{L}(\pi^{\eta}, \eta)$ over $\eta \succeq \mathbf{0}$ (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 $\mathcal{D}(\eta)$ 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 $\delta_k > 0$, ρ , \mathbf{L} , and \mathbf{X} , which generates a state probability $\pi(k)$ ($k = 1, 2, \dots$). This algorithm is described in Algorithm 1 along with the following equations

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

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

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

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), EconCast approximates the difference between the power budget and the average power consumption (22) by observing the dynamics of the energy storage level of 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 η^* as $k \rightarrow \infty$ (see e.g., Theorem 1 of [33]). As mentioned in Section V-F, the choice of parameters σ , δ_k , and τ_k will affect the tradeoff between convergence time and the performance of EconCast.

Part (iii): Convergence Analysis

The detailed proof uses similar techniques as in the proof of Theorem 1 in [33] with minor modifications and can be found in Appendix D.

VII. NUMERICAL RESULTS

In this section, we evaluate the throughput and latency performance of EconCast when operating in groupput and anyput modes. We use the following notation: (i) $\mathcal{T}_g^* (\mathcal{T}_a^*)$ is the oracle groupput (anyput) obtained by solving (P1) or, equivalently, (P2), (ii) $\mathcal{T}_g^\sigma (\mathcal{T}_a^\sigma)$ is the achievable groupput (anyput) of EconCast with a given value of σ obtained by solving (P4), and (iii) $\widetilde{\mathcal{T}}_g^\sigma (\widetilde{\mathcal{T}}_a^\sigma)$ is the groupput (anyput) of EconCast obtained via simulations with a given value of σ . For brevity, we ignore the subscripts of \mathcal{T}^σ when describing results that are general for both groupput and anyput.

A. Setup

We consider clique networks⁵ with $\sigma \in \{0.1, 0.25, 0.5\}$. The nodes' power budgets and consumption levels correspond to the energy harvesting budgets and ultra-low-power transceivers in [7], [8], [34]. 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 $\rho = 10 \mu\text{W}, L = X = 500 \text{ mW}$ 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.

In the simulations, each node has a constant power input at the rate of its power budget, and adjusts the transition rates based on the dynamics of its energy storage level. Since nodes perform carrier sensing when waking up, there are no 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 or the existence of any active listeners, i.e., $\hat{c}(t) = c(t)$ or $\hat{\gamma}(t) = \gamma(t)$.

The simulation results show that \mathcal{T}^σ perfectly matches \mathcal{T}^* for $\sigma \in \{0.25, 0.5\}$. For $\sigma = 0.1$, \mathcal{T}^σ does not converge to \mathcal{T}^* within reasonable time due to the bursty nature of EconCast, as will be described in Section VII-D. 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., $\rho_i = \rho, L_i = L, X_i = X, \forall i \in \mathcal{N}$. The simulation results also confirm that nodes running EconCast consume power on average at the rate of their power budgets.

⁵We evaluate the throughput performance of EconCast in non-clique topologies in VII-E.

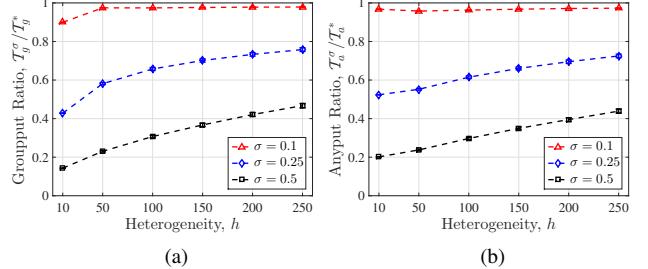


Fig. 2: Sensitivity of the achievable throughput normalized to the oracle throughput, $\mathcal{T}^\sigma/\mathcal{T}^*$, for: (a) groupput and (b) anyput, to the heterogeneity of the power budget, ρ , and power consumption levels, L and X .

B. Heterogeneous Networks – Throughput

One strength of EconCast is its ability to deal with *heterogeneous* networks. Fig. 2 shows the groupput and anyput achieved by EconCast normalized to the corresponding oracle groupput and anyput (i.e., $\mathcal{T}^\sigma/\mathcal{T}^*$), for heterogeneous networks with $N = 5$ and $\sigma \in \{0.1, 0.25, 0.5\}$. 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 independently 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 $500 \mu\text{W}$, and ρ_i has median of $10 \mu\text{W}$ 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. Specifically, in each network sample, each node i samples (ρ_i, L_i, X_i) according to the processes described above. Figs. 2(a) and 2(b) show that the network heterogeneity with respect to both the nodes' power budgets and consumption levels increases as h increases. Fig. 2 also shows that the throughput ratio $\mathcal{T}^\sigma/\mathcal{T}^*$ increases as σ decreases, and approaches 1 as $\sigma \rightarrow 0$. Furthermore, with increased heterogeneity of the network, the throughput ratio has little dependency on the network heterogeneity h but heavy dependency on σ . In general, the groupput and anyput ratios are similar except for homogeneous networks ($h = 10$). In such networks, the anyput ratio is slightly higher than the groupput ratio. This is due to the fact that nodes have the same values of ρ_i, L_i , and X_i). Therefore, determining the existence of any active listeners, $\gamma(t)$, is easier than determining the number of active listeners, $c(t)$.

C. Homogeneous Networks – Throughput and Comparison to Related Work

We now evaluate the throughput of EconCast in *homogeneous* networks with $N = 5$, $\rho = 10 \mu\text{W}$, $L + X = 1 \text{ mW}$, and $\sigma \in \{0.1, 0.25, 0.5\}$. We also compare the groupput achieved by EconCast to three protocols in related work:

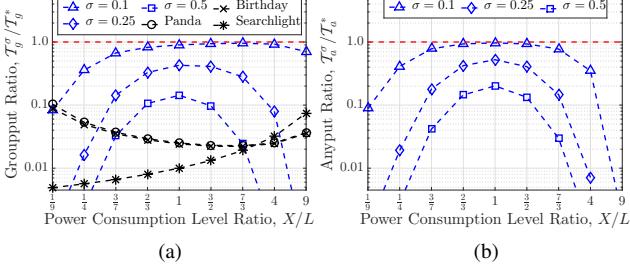


Fig. 3: Throughput performance of EconCast when operating in: (a) groupput mode and (b) anyput mode, with $N = 5$, $\rho = 10 \mu\text{W}$, $L + X = 1 \text{ mW}$, and $\sigma \in \{0.1, 0.25, 0.5\}$, as a function of X/L .

Panda [14], *Birthday* [18], and *Searchlight* [19], 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 N . The throughput of Panda and Birthday is computed as described in [14] and [18], 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 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.

Figs. 3(a) and 3(b) present, respectively, the groupput and anyput achieved by EconCast normalized to the oracle groupput and anyput, as a function of the power consumption ratio X/L , with $N = 5$, $\rho = 10 \mu\text{W}$, and $L + X = 1 \text{ mW}$. Fig. 3(a) also presents the throughput achieved by Panda, Birthday, and Searchlight⁶ protocols. The horizontal dashed lines at 1 represent the oracle groupput and anyput. Note that with $L = X = 500 \mu\text{W}$, the ratio $\mathcal{T}_g^\sigma / \mathcal{T}_g^*$ achieved by EconCast outperforms that of Panda by 6x and 17x with $\sigma = 0.5$ and $\sigma = 0.25$, respectively. In particular, the groupput ratio $\mathcal{T}_g^\sigma / \mathcal{T}_g^*$ significantly outperforms that of prior art for $X \approx L$. The simulation results, which will be discussed later, also verify this throughput improvement.

Fig. 3 shows that for a given value of X/L , \mathcal{T}^σ approaches \mathcal{T}^* with decreasing σ , as expected (see Section V). Moreover, for each value of σ , the throughput ratio $\mathcal{T}^\sigma / \mathcal{T}^*$ increases as the power consumption ratio X/L is closer to 1. This is realistic for current commercial low-power radios that have symmetric power consumption levels in listen and transmit states. This is due to the fact that (i) with small X/L values, nodes enter transmit state infrequently, since listening is expensive and they must pass the listen state to enter the transmit state, and (ii) with large X/L values, nodes spend energy transmitting even when there are no other nodes listening (e.g., $c(t) = 0$). In particular, anyput degrades with

⁶For Searchlight protocol, we compare its throughput upper bound to \mathcal{T}_g^* .

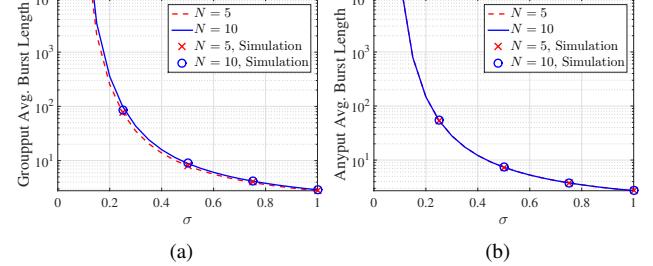


Fig. 4: Analytical (curves) and simulated (markers) average burst length of EconCast when operating in: (a) groupput mode and (b) anyput mode, with $N \in \{5, 10\}$, $\sigma \in \{0.25, 0.5\}$, $\rho = 10 \mu\text{W}$, and $L = X = 500 \mu\text{W}$.

large X/L values, since anyput depends on the *existence* of any active listeners when some node is transmitting. Therefore, when listening is expensive, the fact that multiple nodes listen simultaneously does not improve anyput. 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.

D. Burstiness and Latency

The results until now suggest allowing $\sigma \rightarrow 0$. While reducing σ improves throughput, it considerably increases the *burstiness* of communication, as mentioned in Section V. The burstiness is measured by the *average burst length*, which is defined as the average number of packets that are successfully received in a burst (i.e., the average number of packets a node successfully receives before exiting listen state). The analytical computation of the average burst length can be found in Appendix E. In general, increased burstiness means that the long term throughput can be achieved with given power budgets but the variance of the throughput is more significant during short term intervals.

Figs. 4(a) and 4(b) show the average burst length of EconCast (in log scale) when operating in groupput and anyput modes, respectively, in homogeneous networks with $N \in \{5, 10\}$, $\rho = 10 \mu\text{W}$, $L = X = 500 \mu\text{W}$, and varying σ . Values are obtained using the analytical results (34)–(35) derived in Appendix E (curves) and contrasted with simulations at specific values of σ (markers). Aside from showing that the simulation results and the analytical results are well matched, Fig. 4 also demonstrates that reducing σ dramatically increases burstiness. For example, with $\sigma = 0.25$ and $N = 10$, a node running EconCast in groupput mode has an average burst length of 85, and this value is increased to 4×10^5 with $\sigma = 0.1$. This explains why $\widetilde{\mathcal{T}}_g^\sigma$ does not converge within reasonable time with $\sigma = 0.1$ (see Section VII-A). Comparing Fig. 4(a) with Fig. 4(b), it can be seen that the groupput average burst length increases more rapidly than the anyput average burst length as σ decreases. Moreover, Fig. 4(b) shows that the anyput average burst length is independent of N , which corresponds to the analysis in Appendix E. The reason is that the burst length of EconCast in anyput mode only depends on $\gamma(t)$, which always equals to 1, if the transmission is successful. We remark that reducing the burstiness is a subject of future work.

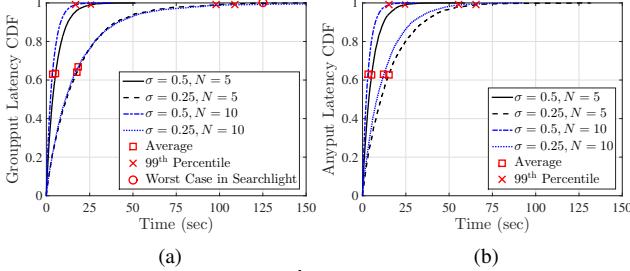


Fig. 5: The CDF, mean, and 99th percentile latency of EconCast when operating in: (a) groupput mode and (b) anyput mode, with $N \in \{5, 10\}$, $\sigma \in \{0.25, 0.5\}$, $\rho = 10 \mu\text{W}$, and $L = X = 500 \mu\text{W}$.

A second metric we consider is the communication *latency*, which 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 more energy, and therefore, need to sleep for longer periods of time. Figs. 5(a) and 5(b) present the CDF latency of EconCast when operating in groupput and anyput modes obtained via simulations, and indicate both the average and the 99th percentile latency values. The homogeneous networks considered are with $N \in \{5, 10\}$, $\sigma \in \{0.25, 0.5\}$, $\rho = 10 \mu\text{W}$, and $L = X = 500 \mu\text{W}$. Fig. 5(a) also shows the pairwise worst case latency of Searchlight computed from [19] under the same power budget and consumption levels.⁷

Fig. 5 shows that the latency increases as σ decreases, since nodes receiving more packets in a short time period (i.e., increased burstiness) have higher variation in their energy storage levels, and need to sleep longer to restore energy. Fig. 5 also shows that a larger value of N results in lower latency, since every node is more likely to receive when more nodes exist in the network. Comparing Fig. 5(a) with Fig. 5(b), it is observed that EconCast operating in anyput mode has slightly lower average latency than in groupput mode. However, with a smaller σ value (i.e., $\sigma = 0.25$), the 99th percentile latency of EconCast when operating in anyput mode is significantly lower than that in groupput mode. This results from the fact that the average burst length of EconCast in anyput mode depends on the existence of any listening nodes, whose value is always less than or equal to the number of listening nodes considered in groupput mode.

For all parameters considered, the 99th percentile groupput 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, its latency is below the worst case latency of Searchlight in most cases (over 99%).

E. Groupput Evaluation in Non-clique Topologies

We now compute the oracle groupput for non-clique topologies (derived in Section IV-C) and evaluate the groupput achieved by EconCast in such scenarios via simulations. Since simultaneous transmissions can happen in non-clique topolo-

⁷This is computed with slot length of 50 ms and a beacon (packet) length of 1 ms as was done in [12].

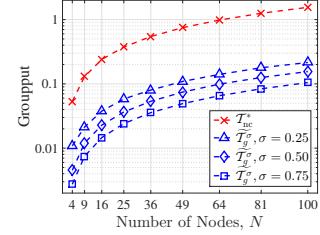


Fig. 6: The oracle groupput, $\mathcal{T}_{\text{nc}}^*$, and the throughput of EconCast when operating in groupput mode obtained via simulations, \mathcal{T}_g^σ , in grid topologies with varying N , $\sigma \in \{0.25, 0.5, 0.75\}$, $\rho = 10 \mu\text{W}$, and $L = X = 500 \mu\text{W}$.

gies, none of the transmissions will be counted as throughput in the simulations.

We use grid topologies with varying number of nodes, N , in which each node has at most 4 neighbors. For example, $N = 25$ represents a 5×5 grid. Fig. 6 presents the oracle groupput, $\mathcal{T}_{\text{nc}}^*$, for grid topologies, and the throughput achieved by EconCast in groupput mode obtained via simulations with varying σ and N . Note that for all the grid topologies considered, the upper and lower bounds of $\mathcal{T}_{\text{nc}}^*$ (see Section IV-C) are the same, providing the exact oracle groupput.

Fig. 6 shows that EconCast achieves 14% – 22% of the oracle groupput, $\mathcal{T}_{\text{nc}}^*$, with $\sigma = 0.25$. Although increasing σ leads to lower groupput, it can be observed that as N increases, the groupput approaches 10% of $\mathcal{T}_{\text{nc}}^*$ with $\sigma = 0.5$. Despite the fact that the groupput cannot be obtained for $\sigma = 0.1$, achieving 10% – 20% of $\mathcal{T}_{\text{nc}}^*$ is remarkable given the fact that EconCast operates in a distributed manner.

VIII. EXPERIMENTAL EVALUATION

To experimentally evaluate the performance of EconCast-C,⁸ we implement it using the Texas Instruments eZ430-RF2500-SEH node [35].⁹ In this section, we first describe the energy measurements performed on the nodes running EconCast-C. Then, we describe the method by which nodes can estimate the number of listening nodes. Finally, we experimentally evaluate the performance of EconCast-C.

A. 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.¹⁰ The power consumption levels are very similar from node to node. Recall from Section VII that the performance of EconCast

⁸See Section V-D the reasons for only implementing EconCast-C.

⁹A demonstration of the testbed is presented in [36].

¹⁰This corresponds to a -16 dBm transmission power, at which nodes within the same room typically have little or no packet loss.

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 $\rho \in \{10 \mu\text{W}, 50 \mu\text{W}\}$, $L = 670 \mu\text{W}$, and $X = 560 \mu\text{W}$).

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, 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.

B. 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 ($C_{\text{cap}} = 5 \text{ F}$) that stores energy in advance. The energy consumed is computed by

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

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

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

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,¹¹ 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 $\bar{\mathcal{T}}_g^\sigma$ and \mathcal{T}_g^σ , respectively. In Section VIII-D, we compare the experimental throughput to both $\bar{\mathcal{T}}_g^\sigma$ and \mathcal{T}_g^σ . Having verified the power consumption of the nodes, we replace the capacitor with AAA batteries,¹² allowing the experiments to run for longer times.

¹¹The additional energy usage includes the energy consumed in powering up the regulator circuitry, etc.

¹²The constant voltage of AAA batteries limits the ability to measure the power consumption of the nodes.

C. 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, $\lambda_{xl}^C(t)$, 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 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.

D. Performance Evaluation

We consider homogeneous networks with $N \in \{5, 10\}$, $\rho \in \{1 \text{ mW}, 5 \text{ mW}\}$, $\sigma \in \{0.25, 0.5\}$, and nodes are 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: Fig. 7 presents the ratio of the experimentally obtained throughput, $\bar{\mathcal{T}}_g^\sigma$, normalized to the achievable throughput \mathcal{T}_g^σ and $\bar{\mathcal{T}}_g^\sigma$ (see Section VIII-B). Separate charts represent the results for differing number of nodes, N , and power budget, ρ . Points marked “Ideal” show the experimental throughput normalized to the achievable throughput computed by solving (P4) with the target power budget ρ (i.e., $\bar{\mathcal{T}}_g^\sigma / \mathcal{T}_g^\sigma$). 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., $\bar{\mathcal{T}}_g^\sigma / \bar{\mathcal{T}}_g^\sigma$). As expected, $\bar{\mathcal{T}}_g^\sigma$ is higher than \mathcal{T}_g^σ , resulting in a lower throughput ratio.

Fig. 7 shows that despite the practical limitations (e.g., packet collisions and inaccurate clocks) faced when running EconCast-C on real hardware, the ratio $\bar{\mathcal{T}}_g^\sigma / \bar{\mathcal{T}}_g^\sigma$ is between 57% – 77% ($\bar{\mathcal{T}}_g^\sigma / \mathcal{T}_g^\sigma$ is between 67% – 81%) for all settings considered. Moreover, Table III shows the improvement of EconCast-C over the throughput of Panda computed according to [14], 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.

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 VIII-C) after each packet transmission which effectively reduces the number of bits delivered. Second, collisions of pings or failed decodings of

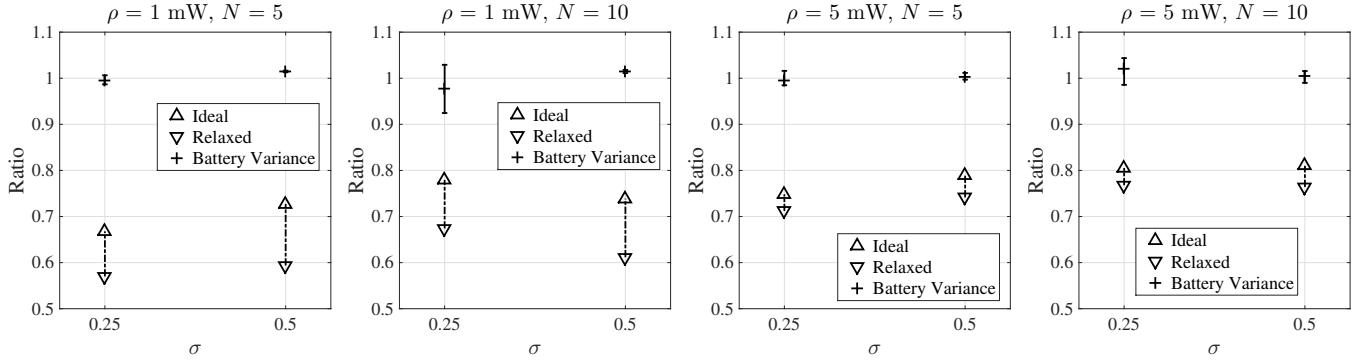


Fig. 7: Points marked as “Ideal” (“Relaxed”) represent ratio of experimental throughput normalized to the achievable throughput obtained by using the target power budget (actual power consumption) and points marked as “Battery Variance” present the average, minimum, and maximum ratios of power consumption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

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

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

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: In Section VIII-B, we show that the power consumption of the virtual battery is valid for evaluating the actual power consumption of the node. Fig. 7 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 VII-D), 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 transmitter 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 IV presents the distribution of number of pings (i.e., 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

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

# 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

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 Fig. 7.

IX. 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 also provided efficient methods to obtain oracle groupput and oracle anyput for a given set of heterogeneous nodes.

We developed the EconCast 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 in the network, and power budgets and consumption levels) can achieve the oracle groupput and anyput in a limiting sense (when $\sigma \rightarrow 0$).

We evaluated the throughput performance of EconCast numerically and through extensive simulations, and compared it to the state of the art. We also considered the design tradeoffs in relation to σ and the impact of σ on the burstiness and throughput. Finally, we experimentally evaluated EconCast using commercial-off-the-shelf energy harvesting 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., [4]), that have improved energy awareness compared to the TI eZ430-RF2500-SEH nodes, would enable to better assess the tradeoffs related to the protocol design. Finally, considering unique application characteristics and their relation to groupput and anyput is an open problem.

REFERENCES

[1] T. Chen, J. Ghaderi, D. Rubenstein, and G. Zussman, "Maximizing broadcast throughput under ultra-low-power constraints," in *Proc. ACM CoNEXT'16*, 2016.

[2] L. Atzori, A. Iera, and G. Morabito, "The internet of things: A survey," *Computer Networks*, vol. 54, no. 15, pp. 2787 – 2805, 2010.

[3] M. Gorlatova, P. Kinget, I. Kymmissis, D. Rubenstein, X. Wang, and G. Zussman, "Challenge: ultra-low-power energy-harvesting active networked tags (EnHANTs)," in *Proc. ACM MobiCom'09*, 2009.

[4] R. Margolies, M. Gorlatova, J. Sarik, G. Stajic, J. Zhu, P. Miller, M. Szczerzak, B. Vigraham, L. Carloni, P. Kinget *et al.*, "Energy-harvesting active networked tags (EnHANTs): Prototyping and experimentation," *ACM Trans. Sens. Netw.*, vol. 11, no. 4, p. 62, 2015.

[5] M. Buettner, B. Greenstein, and D. Wetherall, "Dewdrop: an energy-aware runtime for computational RFID," in *Proc. USENIX NSDI'11*, 2011.

[6] 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.

[7] M. Gorlatova, A. Wallwater, and G. Zussman, "Networking low-power energy harvesting devices: Measurements and algorithms," *IEEE Trans. Mobile Comput.*, vol. 12, no. 9, pp. 1853–1865, Sept. 2013.

[8] M. Gorlatova, J. Sarik, G. Grebla, M. Cong, I. Kymmissis, and G. Zussman, "Movers and shakers: Kinetic energy harvesting for the internet of things," *IEEE J. Sel. Areas Commun.*, vol. 33, no. 8, pp. 1624–1639, 2015.

[9] 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.*, vol. 33, no. 3, pp. 360–381, Mar. 2015.

[10] "Texas Instruments CC2640 SimpleLink Bluetooth smart wireless MCU," 2015. [Online]. Available: <http://ti.com/lit/ds/symlink/cc2640.pdf>

[11] 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.

[12] W. Sun, Z. Yang, K. Wang, and Y. Liu, "Hello: A generic flexible protocol for neighbor discovery," in *Proc. IEEE INFOCOM'14*, 2014.

[13] 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.

[14] R. Margolies, G. Grebla, T. Chen, D. Rubenstein, and G. Zussman, "Panda: Neighbor discovery on a power harvesting budget," *IEEE J. Sel. Areas Commun.*, vol. 34, no. 12, pp. 3606–3619, Dec. 2016.

[15] 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.

[16] T. Sriskoosai, K. Keamarungsi, P. Lamsrichan, and K. Araki, "Practical data compression in wireless sensor networks: A survey," *J. Netw. Comput. Appl.*, vol. 35, no. 1, pp. 37–59, 2012.

[17] 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.

[18] 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.

[19] M. Bakht, M. Trower, and R. H. Kravets, "Searchlight: Won't you be my neighbor?" in *Proc. ACM MobiCom'12*, 2012.

[20] 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.*, vol. 16, no. 3, pp. 1448–1459, 2014.

[21] W. Ye, J. Heidemann, and D. Estrin, "Medium access control with coordinated adaptive sleeping for wireless sensor networks," *IEEE/ACM Trans. Netw.*, vol. 12, no. 3, pp. 493–506, 2004.

[22] P. Dutta and D. Culler, "Practical asynchronous neighbor discovery and rendezvous for mobile sensing applications," in *Proc. ACM SenSys'08*, 2008.

[23] 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.

[24] 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.

[25] S. Vasudevan, M. Adler, D. Goeckel, and D. Towsley, "Efficient algorithms for neighbor discovery in wireless networks," *IEEE/ACM Trans. Netw.*, vol. 21, no. 1, pp. 69–83, 2013.

[26] M. Chiang, S. H. Low, J. C. Doyle *et al.*, "Layering as optimization decomposition: A mathematical theory of network architectures," *Proc. IEEE*, vol. 95, no. 1, pp. 255–312, 2007.

[27] J. Ghaderi and R. Srikant, "On the design of efficient CSMA algorithms for wireless networks," in *Proc. IEEE CDC'10*, 2010.

[28] J. Ghaderi, S. Borst, and P. Whiting, "Queue-based random-access algorithms: Fluid limits and stability issues," *Stochastic Systems*, vol. 4, no. 1, pp. 81–156, 2014.

[29] P. Marbach and A. Eryilmaz, "A backlog-based CSMA mechanism to achieve fairness and throughput-optimality in multi-hop wireless networks," in *Proc. Allerton Conf.*, 2008.

[30] L. Jiang and J. Walrand, "A distributed CSMA algorithm for throughput and utility maximization in wireless networks," *IEEE/ACM Trans. Netw.*, vol. 18, pp. 960–972, 2010.

[31] K. Xu, O. Dousse, and P. Thiran, "Self-synchronizing properties of CSMA wireless multi-hop networks," in *Proc. ACM SIGMETRICS'10*, 2010.

[32] 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.*, vol. 22, no. 4, pp. 1325–1336, 2014.

[33] L. Jiang and J. Walrand, "Convergence and stability of a distributed CSMA algorithm for maximal network throughput," in *Proc. CDC'09*, 2009.

[34] K. Philips, "Ultra low power short range radios: Covering the last mile of the IoT," in *Proc. IEEE ESSCIRC'14*, 2014.

[35] "Texas Instruments eZ430-RF2500-SEH solar energy harvesting development tool user's guide," 2013.

[36] 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.

[37] P. Diaconis and D. Stroock, "Geometric bounds for eigenvalues of markov chains," *Ann. Appl. Probab.*, vol. 1, pp. 36–61, 1991.

APPENDIX A

PROOF OF LEMMA 1

We describe here a schedule that works assuming each packet transmission has a fixed transmit length of θ , though the proof can be extended to varying transmission lengths. Therefore, time can be broken into slots of length θ (by the oracle) and nodes sleep, listen, or transmit on a per-slot basis.

Assume that the optimal solution (α^*, β^*) to $(P2)$ yields rational values for all α_i^* and β_i^* . The period size of the oracle schedule P is set to the least common denominator over all solution variables in (α^*, β^*) . Hence, during each period, node i listens for α_i^*P slots and transmits for β_i^*P slots. The slots during the period can be assigned arbitrarily by the oracle to transmitters (e.g., weighted round-robin or in-order) and (11) ensures that there will be sufficient slots. Once the slots for the transmitters are assigned, each listener can then choose their α_i^*P slots in which they listen from the set of transmit slots assigned to other transmitters, and (12) ensures these are sufficient as well (note that multiple listeners are permitted for a single transmitter slot).

If the periodic scheduler is launched immediately, some nodes may not have the harvested (or budgeted) energy to perform all listen and transmit tasks within the first period (e.g., it may be assigned to transmit and listen early on, and recoup the energy during later slots). If such a case, we simply delay the initial iteration, allowing all nodes to harvest (or budget) $\rho_i P\theta$ amount of energy for that initial period. Then the nodes have enough energy to repeat all the subsequent periods since the energy node i spends during the k -th period is $(\alpha_i L_i P\theta + \beta_i X_i P\theta)$, and the energy it accumulates during this period is $\rho_i P\theta$. Hence, it follows from (9) that no more

energy is spent than is accumulated (budgeted). Therefore, there is sufficient energy to repeat the period. \blacksquare

APPENDIX B SOLVING (P2) AND (P3)

The optimization problem (P2) can be solved via two steps. First, we show that in the optimal solution to (P2), the equalities strictly hold in constraints (9) and (12). Next, we solve for the optimal solution. Note that in homogeneous networks, the constraints become $\alpha L + \beta X \leq \rho$ and $\alpha \leq (N-1)\beta$, and $\alpha_i^* = \alpha^*$ and $\beta_i^* = \beta$.

We prove the first part by contradiction. Note that if both inequalities strictly hold, α can always be increased, resulting in higher throughput. Therefore, it suffices that at least one of the two inequalities is satisfied with strict equality.

Case 1: If $\alpha^* L + \beta^* X = \rho - \epsilon$ for some $\epsilon > 0$ and $\alpha^* = (N-1)\beta^*$, the optimal solution is given by $\beta^* = \frac{\rho - \epsilon}{X + (N-1)L}$, $\alpha^* = (N-1)\beta^*$. Let a new solution be

$$\alpha' = \alpha^* + \frac{(N-1)\epsilon}{X + (N-1)L}, \quad \beta' = \beta^* + \frac{\epsilon}{X + (N-1)L},$$

and it can be verified that (α', β') satisfies: (i) $\alpha' L + \beta' X = \rho$ and $\alpha' = (N-1)\beta'$, and (ii) $(\mathcal{T}^*)' = N\alpha' > \mathcal{T}^*$.

Case 2: If $\alpha^* L + \beta^* X = \rho$ and $\alpha^* = (N-1)\beta^* - \delta$ for some $\delta > 0$, the optimal solution is given by $\beta^* = \frac{\rho + \delta L}{X + (N-1)L}$, $\alpha^* = (N-1)\beta^* - \delta$. Let a new solution be

$$\alpha'' = \alpha^* + \frac{\delta X}{X + (N-1)L}, \quad \beta'' = \beta^* - \frac{\delta L}{X + (N-1)L},$$

and it can be verified that (α'', β'') satisfies: (i) $\alpha'' L + \beta'' X = \rho$ and $\alpha'' = (N-1)\beta''$, (ii) $\beta'' > 0$ still holds, and (iii) $(\mathcal{T}^*)'' = N\alpha'' > \mathcal{T}^*$.

Next, given that equalities strictly hold in constraints (9) and (12), the optimal solution of (α^*, β^*) can be obtained by solving $\alpha^* L + \beta^* X = \rho$ and $\alpha^* = (N-1)\beta^*$, and the solution is described in Section IV-A. Similarly, (P3) can be solved to obtain the results described in Section IV-B. \blacksquare

APPENDIX C PROOF OF LEMMA 2

For the capture version EconCast-C, we prove that the transition rates (18a), (18b), (18c), and (18e) will drive the network Markov chain to a steady state with distribution (19), by checking that the *detailed balanced equations* hold. We assume $\sigma = 1$ and drop the constant term $1/\sigma$ for brevity. For network state \mathbf{w} , define $\mathcal{N}_s(\mathbf{w})$, $\mathcal{N}_l(\mathbf{w})$, and $\mathcal{N}_x(\mathbf{w})$ as the sets of nodes in sleep, listen, and transmit states, respectively, and their cardinalities as $N_s(\mathbf{w})$, $N_l(\mathbf{w})$, and $N_x(\mathbf{w})$, respectively. Note that $N_s(\mathbf{w}) + N_l(\mathbf{w}) + N_x(\mathbf{w}) = N$ and $N_x(\mathbf{w}) \in \{0, 1\}$. We also use $\mathbf{w} = (\mathcal{N}_s(\mathbf{w}), \mathcal{N}_l(\mathbf{w}), \mathcal{N}_x(\mathbf{w}))$ to denote network state \mathbf{w} and $r(\mathbf{w}, \mathbf{w}')$ to denote the transition rate from state \mathbf{w} to \mathbf{w}' . We consider the following cases.

Case 1: If node i is in sleep state ($w_i = s$), the only transition that can happen is to transition into listen state when the channel is clear, i.e., $\mathbf{w} = (\mathcal{N}_s(\mathbf{w}), \mathcal{N}_l(\mathbf{w}), \emptyset) \rightarrow \mathbf{w}' = (\mathcal{N}_s(\mathbf{w}) \setminus \{i\}, \mathcal{N}_l(\mathbf{w}) \cup \{i\}, \emptyset)$. In this case, $r(\mathbf{w}, \mathbf{w}') = \exp(-\eta_i L_i)$ and $r(\mathbf{w}', \mathbf{w}) = 1$.

Case 2: If node i is in listen state ($w_i = l$), and transitions to sleep state, i.e., $\mathbf{w} = (\mathcal{N}_s(\mathbf{w}), \mathcal{N}_l(\mathbf{w}), \emptyset) \rightarrow \mathbf{w}' =$

$(\mathcal{N}_s(\mathbf{w}) \cup \{i\}, \mathcal{N}_l(\mathbf{w}) \setminus \{i\}, \emptyset)$. In this case, $r(\mathbf{w}, \mathbf{w}') = 1$ and $r(\mathbf{w}', \mathbf{w}) = \exp(-\eta_i L_i)$.

Case 3: If node i is in listen state ($w_i = l$) and transitions to transmit state, i.e., $\mathbf{w} = (\mathcal{N}_s(\mathbf{w}), \mathcal{N}_l(\mathbf{w}), \emptyset) \rightarrow \mathbf{w}' = (\mathcal{N}_s(\mathbf{w}), \mathcal{N}_l(\mathbf{w}) \setminus \{i\}, \{i\})$. In this case, $r(\mathbf{w}, \mathbf{w}') = \exp(\eta_i(L_i - X_i))$ and $r(\mathbf{w}', \mathbf{w}) = \exp(-N_l(\mathbf{w}'))$.

Case 4: If node i is in transmit state ($w_i = x$), the only transition that can happen is to transition to listen state when its transmission is finished, i.e., $\mathbf{w} = (\mathcal{N}_s(\mathbf{w}), \mathcal{N}_l(\mathbf{w}), \{i\}) \rightarrow \mathbf{w}' = (\mathcal{N}_s(\mathbf{w}) \setminus \{i\}, \mathcal{N}_l(\mathbf{w}) \cup \{i\}, \emptyset)$. In this case, $r(\mathbf{w}, \mathbf{w}') = \exp(-N_l(\mathbf{w}))$ and $r(\mathbf{w}', \mathbf{w}) = \exp(\eta_i(L_i - X_i))$.

For each case, $\pi_{\mathbf{w}} \cdot r(\mathbf{w}, \mathbf{w}') = \pi_{\mathbf{w}'} \cdot r(\mathbf{w}', \mathbf{w})$ holds and similar detailed balance equations hold for the non-capture version EconCast-NC. Therefore, we complete the proof. \blacksquare

APPENDIX D PROOF OF PART (III)

We use the same notation described in Appendix C. In addition, denote $W = |\mathcal{W}|$ as the number of network states. Let $t_0 = 0$ and recall that the length of the k -th interval is $\tau_k = t_k - t_{k-1}$. At time t , the probability of the system being in state \mathbf{w} is denoted by $\zeta_{\mathbf{w}}(t)$. Notice that throughout the k -th interval, the vector of Lagrange multipliers $\boldsymbol{\eta}(k-1)$ remains unchanged and is updated at time t_k according to

$$\eta_i(k) = [\eta_i(k-1) - \delta_k(\rho_i - \hat{\mu}_i(k))]^+, \quad \forall i \in \mathcal{N} \quad (27)$$

in which $\hat{\mu}_i(k)$ is the *empirical* energy consumption rate of node i in the k -th interval.¹³ The relationship between the empirical energy consumption rate and the energy storage level $b_i(t)$ of node i is given by $\rho_i - \hat{\mu}_i(k) = [b_i(t_k) - b_i(t_{k-1})]/\tau_k$.

In addition, let $\mathbf{w}^0(k)$ denote the state of the network Markov chain at the beginning of the k -th interval (i.e., at time t_{k-1}). Define the random vector $\mathbf{U}(k-1) := [\hat{\mu}(k-1), \boldsymbol{\eta}(k-1), \mathbf{w}^0(k-1)]$. For $k \geq 1$, let \mathcal{F}_{k-1} be the σ -field generated by $\mathbf{U}(0), \mathbf{U}(1), \dots, \mathbf{U}(k-1)$.

In the k -th interval, define the gradient vector as $\mathbf{g}(k) = [g_i(k)]$, in which $g_i(k) = \rho_i(k) - \mu_i(k)$. Given vector $\boldsymbol{\eta}(k-1)$, $\mathbf{g}(k)$ is a gradient of $\mathcal{L}(\boldsymbol{\eta})$ (recall that the dual problem of (P4) is $\min_{\boldsymbol{\eta} \geq 0} \mathcal{L}(\boldsymbol{\eta})$). However, EconCast follows (27) and only has an empirical estimation $\hat{\mathbf{g}}(k) = [\hat{g}_i(k)]$ of $\mathbf{g}(k)$, in which $\hat{g}_i(k) = \rho_i - \hat{\mu}_i(k)$. The “error” term is given by $\mathbf{g}^{\text{err}}(k) = [g_i^{\text{err}}(k)]$, in which

$$g_i^{\text{err}}(k) = \mathbb{E}[\hat{g}_i(k)|\mathcal{F}_{k-1}] - g_i(k) = \mu_i(k) - \mathbb{E}[\hat{\mu}_i(k)|\mathcal{F}_{k-1}]. \quad (28)$$

The “noise” term is given by $\mathbf{g}^{\text{noise}}(k) = [g_i^{\text{noise}}(k)]$, in which

$$g_i^{\text{noise}}(k) = \mathbb{E}[\hat{\mu}_i(k)|\mathcal{F}_{k-1}] - \hat{\mu}_i(k). \quad (29)$$

Therefore we can write the empirical gradient as

$$\hat{g}_i(k) = g_i(k) + g_i^{\text{err}}(k) + g_i^{\text{noise}}(k), \quad \forall i \in \mathcal{N}.$$

Notice that since $\rho_i(k)$ and $\hat{\mu}_i(k)$ are both bounded, the noise term is also bounded by some constant, i.e., $|g_i^{\text{noise}}(k)| \leq g_{\max}^{\text{noise}}, \forall i \in \mathcal{N}$.

Below, we show that with EconCast, the error term $\mathbf{g}^{\text{err}}(k)$ (28) decreases “fast enough” with time. We focus on the $(k+1)$ -th interval and denote the Continuous-Time Markov chain (CTMC) in the $(k+1)$ -th interval by

¹³Although we assume a constant power budget of ρ_i at each node i , the proof can be easily extended to scenarios where ρ_i varies with time.

$X(t)$, $\forall t \in [t_k, t_{k+1}]$, whose transition rate matrix is denoted as $\mathbf{Q} = [Q(\mathbf{w}, \mathbf{w}')$. For brevity, we assume $\sigma = 1$ and drop the term $1/\sigma$. Recall from (19), $Z^{\eta(k)}$ is a normalizing constant which can be easily bounded by

$$Z^{\eta(k)} \leq W \cdot \exp(N) = (2+N) \cdot 2^{N-1} \cdot \exp(N),$$

in which we use the fact that $\mathcal{T}_{\mathbf{w}} \leq N$. On the other hand, denote $\bar{C} = \max_{i \in \mathcal{N}} \{\max\{L_i, X_i\}\}$, we also have

$$\pi_{\mathbf{w}}^{\eta(k)} \cdot Z^{\eta(k)} \geq \exp[-N \cdot \bar{C} \cdot \max_i \{\eta_i(k)\}].$$

From (27), $\max_i \{\eta_i(k)\}$ can be further bounded by $\max_{i \in \mathcal{N}} \{\eta_i(k)\} \leq \bar{C} \cdot \sum_{m=1}^k \delta_m$. Therefore, denoting $\xi_k := \bar{C}^2 \cdot \sum_{m=1}^k \delta_m$, we have the minimum probability in the stationary distribution lower bounded by

$$\pi_{\mathbf{w}}^{\eta(k)} \geq \exp(-\xi_k N) (W \cdot \exp(N))^{-1} := \pi_{\min}^{\eta(k)}. \quad (30)$$

Next, we state the following useful proposition.

Proposition 1: For the Markov chain $X(t)$ in the $(k+1)$ -th interval, if $Q(\mathbf{w}, \mathbf{w}') > 0$, then there exists $Q_{k+1}^{\max} \geq Q_{k+1}^{\min} > 0$ such that $Q_{k+1}^{\min} \leq Q(\mathbf{w}, \mathbf{w}') \leq Q_{k+1}^{\max}$.

Proof: Given network state \mathbf{w} , there is $N_s(\mathbf{w}) + 2N_l(\mathbf{w}) + N_x(\mathbf{w}) \leq 2N$ states \mathbf{w}' other than \mathbf{w} that $X(t)$ can transition to. For any state $\mathbf{w}' \neq \mathbf{w}$, we have for each node i ,

$$Q(\mathbf{w}, \mathbf{w}') = \begin{cases} \exp[-\eta_i(k)L_i], & w_i : s \rightarrow l, \\ \exp[\eta_i(k)(L_i - X_i)], & w_i : l \rightarrow x, \\ 1, & w_i : l \rightarrow s, \\ \exp(-\mathcal{T}_{\mathbf{w}}), & w_i : x \rightarrow l. \end{cases} \quad (31)$$

It is easy to see that $Q_{k+1}^{\min} = \min\{\exp(-N), \exp(-\xi_k)\}$ suffices. In particular, given the network size N , for sufficiently large k , $\xi_k > N$ holds, and therefore we use $Q_{k+1}^{\min} := \exp(-\xi_k)$.

On the other hand, if for $L_i \leq X_i, \forall i \in \mathcal{N}$, each one of the $2N$ transition rates in (31) is less than or equal to 1, and $Q_{k+1}^{\max} = 2N$ is sufficient. Second, if there exists some node i such that $L_i > X_i$, because of the costs are always upper bounded by \bar{C} , the transition rate is also upper bounded by $\exp(\bar{C}\eta_{\max}) \leq \exp(\bar{C}^2 \sum_{m=1}^k \delta_m) = \exp(\xi_k)$. Therefore $Q_{k+1}^{\max} = 2N \exp(\xi_k)$ is sufficient. ■

Following the standard method, we perform uniformization on $X(t)$ whose transition rate matrix is denoted as \mathbf{Q} . If each element of \mathbf{Q} has an absolute value less than or equal to 1, then we can write $X(t) = Y(M(t))$, in which $Y(n)$ is a discrete time Markov chain with probability transition matrix $\mathbf{P} = \mathbf{I} + \mathbf{Q}/Q_{k+1}^{\max}$ and \mathbf{I} is the identity matrix. $M(t)$ is an independent Poisson process with rate Q_{k+1}^{\max} .

Then, we estimate how far the empirical power consumption $\mathbb{E}[\hat{\mu}_i(k+1)|\mathcal{F}_k]$ is away from the desired value $\mu_i(k+1)$ (under fixed $\eta(k)$). Denote the probability of $X(t)$ in state \mathbf{w} at time $t \in [t_k, t_{k+1}]$ by $\zeta_{\mathbf{w}}(t)$ and $\zeta(t) = [\zeta_{\mathbf{w}}(t)]$. Given the initial state at time t_k is $\mathbf{w}^0(k+1)$, we have

$$\begin{aligned} & \mathbb{E}[\hat{\mu}_i(k+1)|\mathcal{F}_k] \\ &= \mathbb{E} \left[\int_{t_k}^{t_{k+1}} \sum_{\mathbf{w} \in \mathcal{W}} (\mathbf{1}_{\{w_i=l\}} L_i + \mathbf{1}_{\{w_i=x\}} X_i) dt / \tau_{k+1} \right] \\ &= \int_{t_k}^{t_{k+1}} \sum_{\mathbf{w} \in \mathcal{W}} (\mathbb{P}\{w_i(t) = l\} L_i + \mathbb{P}\{w_i(t) = x\} X_i) dt / \tau_{k+1} \\ &= \frac{L_i}{\tau_{k+1}} \sum_{\mathbf{w} \in \mathcal{W}_i^l} \int_{t_k}^{t_{k+1}} \zeta_{\mathbf{w}}(t) dt + \frac{X_i}{\tau_{k+1}} \sum_{\mathbf{w} \in \mathcal{W}_i^x} \int_{t_k}^{t_{k+1}} \zeta_{\mathbf{w}}(t) dt \\ &= L_i \cdot \hat{\alpha}_i(k+1) + X_i \cdot \hat{\beta}_i(k+1), \end{aligned}$$

in which

$$\hat{\alpha}_i(k+1) = \sum_{\mathbf{w} \in \mathcal{W}_i^l} \int_{t_k}^{t_{k+1}} \zeta_{\mathbf{w}}(t) dt / \tau_{k+1},$$

$$\hat{\beta}_i(k+1) = \sum_{\mathbf{w} \in \mathcal{W}_i^x} \int_{t_k}^{t_{k+1}} \zeta_{\mathbf{w}}(t) dt / \tau_{k+1},$$

are the empirical average probabilities that node i spends in listen and transmit states in the $(k+1)$ -th interval.

Let θ_2 be the Second Largest Eigenvalue modulus of the transition probability matrix $\mathbf{P} = \mathbf{I} + \mathbf{Q}/Q_{k+1}^{\max}$. By Theorem [], this total variation distance can be bounded by

$$\begin{aligned} \|\zeta(t) - \pi^{\eta(k)}\|_{\text{TV}} &= \frac{1}{2} \sum_{\mathbf{w} \in \mathcal{W}} |\zeta_{\mathbf{w}}(t) - \pi_{\mathbf{w}}^{\eta(k)}| \\ &\leq \frac{1}{2} \sqrt{\frac{1 - \pi_{\mathbf{w}^0}^{\eta(k)}}{\pi_{\mathbf{w}^0}^{\eta(k)}}} \exp(-Q_{k+1}^{\max}(1 - \theta_2)t) \\ &\leq \frac{1}{2} \sqrt{\frac{1}{\pi_{\min}^{\eta(k)}}} \exp(-Q_{k+1}^{\max}(1 - \theta_2)t). \end{aligned} \quad (32)$$

Also, the Second Largest Eigenvalue, θ_2 , can be bounded by Cheeger's inequality [37], i.e.,

$$\theta_2 \leq 1/\phi^2/2 \Leftrightarrow 1/(1 - \theta_2) \leq 2/\phi^2,$$

where ϕ is the “conductance” of \mathbf{P} , which satisfies the following inequality

$$\phi \geq \min_{\mathbf{w} \neq \mathbf{w}', P(\mathbf{w}, \mathbf{w}') > 0} \{\pi_{\min}^{\eta(k)} P(\mathbf{w}, \mathbf{w}')\} \geq \pi_{\min}^{\eta(k)} \cdot Q_{k+1}^{\min} / Q_{k+1}^{\max}.$$

Therefore we obtain

$$1/(1 - \theta_2) \leq 2/\phi^2 \leq 2(Q_{k+1}^{\max})^2 / (\pi_{\min}^{\eta(k)} \cdot Q_{k+1}^{\min})^2. \quad (33)$$

Putting everything into the estimation, we have

$$\begin{aligned} & |\mathbb{E}[\hat{\mu}_i(k+1)|\mathcal{F}_k] - \mu_i(k+1)| \\ &= |L_i(\hat{\alpha}_i(k+1) - \alpha_i(k+1)) + X_i(\hat{\beta}_i(k+1) - \beta_i(k+1))| \\ &\leq L_i|\hat{\alpha}_i(k+1) - \alpha_i(k+1)| + X_i|\hat{\beta}_i(k+1) - \beta_i(k+1)|. \end{aligned}$$

Furthermore, we have

$$\begin{aligned} & |\hat{\alpha}_i(k+1) - \alpha_i(k+1)| \\ &= \left| \sum_{\mathbf{w} \in \mathcal{W}_i^l} \int_{t_k}^{t_{k+1}} (\zeta_{\mathbf{w}}(t) - \pi_{\mathbf{w}}^{\eta(k)}) dt / \tau_{k+1} \right| \\ &\leq \int_{t_k}^{t_{k+1}} \sum_{\mathbf{w} \in \mathcal{W}_i^l} |\zeta_{\mathbf{w}}(t) - \pi_{\mathbf{w}}^{\eta(k)}| dt / \tau_{k+1}. \end{aligned}$$

Similarly, $|\hat{\beta}_i(k+1) - \beta_i(k+1)| \leq \int_{t_k}^{t_{k+1}} \sum_{\mathbf{w} \in \mathcal{W}_i^x} |\zeta_{\mathbf{w}}(t) - \pi_{\mathbf{w}}^{\eta(k)}| dt / \tau_{k+1}$. Adding both terms together yields

$$\begin{aligned} & L_i|\hat{\alpha}_i(k+1) - \alpha_i(k+1)| + X_i|\hat{\beta}_i(k+1) - \beta_i(k+1)| \\ &\leq \bar{C} \int_{t_k}^{t_{k+1}} \sum_{\mathbf{w} \in \mathcal{W}} |\zeta_{\mathbf{w}}(t) - \pi_{\mathbf{w}}^{\eta(k)}| dt / \tau_{k+1} \\ &\leq \frac{\bar{C}}{\sqrt{\pi_{\min}^{\eta(k)}}} \int_{t_k}^{t_{k+1}} \exp(-Q_{k+1}^{\max}(1 - \theta_2)t) dt / \tau_{k+1} \quad (\text{use (32)}) \\ &\leq \frac{\bar{C}}{\sqrt{\pi_{\min}^{\eta(k)}}} \int_0^{t_{k+1}} \exp(-Q_{k+1}^{\max}(1 - \theta_2)t) dt / \tau_{k+1} \\ &\leq \frac{\bar{C}}{\sqrt{\pi_{\min}^{\eta(k)}}} \cdot [2Q_{k+1}^{\max}(1 - \theta_2)]^{-1} / \tau_{k+1} \\ &\leq \frac{\bar{C}Q_{k+1}^{\max}}{\tau_{k+1}(Q_{k+1}^{\min})^2(\pi_{\min}^{\eta(k)})^{5/2}} := f(k) / \tau_{k+1}, \quad (\text{use (33)}) \end{aligned}$$

in which $f(k) := \frac{\bar{C}Q_{k+1}^{\max}}{(Q_{k+1}^{\min})^2(\pi_{\min}^{\eta(k)})^{5/2}}$. Therefore we have,

$$|\mathbb{E}[\hat{\mu}_i(k+1)|\mathcal{F}_k] - \mu_i(\eta(k))| \leq f(k) / \tau_{k+1},$$

and the error term, defined in (28), satisfies $|g_i^{\text{err}}(k+1)| \leq f(k) / \tau_{k+1}$. Therefore, with defined $f(k)$ and proper choice of τ_{k+1} (e.g., $\tau_{k+1} = k+1$), the error term is diminishing.

Based on this result and following similar steps in the proof of Theorem 1 in [33], it can be shown that η converges to η^*

with probability 1. Therefore, we complete the proof. ■

APPENDIX E BURSTINESS ANALYSIS OF ECONCAST

To derive the average burst length (denoted by B) of EconCast-C, we use π_w^* to denote the optimal solution to (P4) and define $\mathcal{W}' = \{w \in \mathcal{W} : \nu_w = 1, c_w \geq 1\}$, i.e., the set of states with successfully received bursts. Recall that for a given value of σ , the optimal value of (P4) is exactly \mathcal{T}^σ . According to (18e), for a given state $w \in \mathcal{W}'$, the average burst length of EconCast-C in groupput mode is $\exp(c_w/\sigma)$. Therefore, during a (long enough) time duration of T , the average number of bursts received by all the nodes can be computed by $\sum_{w \in \mathcal{W}'} \frac{T \cdot \pi_w^*}{\exp(c_w/\sigma)}$, and the average burst length

of EconCast in groupput mode, B_g , is given by

$$B_g = (\text{Avg. Total Burst Length}) / (\text{Avg. Number of Bursts}) \\ = \frac{T \cdot \sum_{w \in \mathcal{W}'} \pi_w^*}{\sum_{w \in \mathcal{W}'} \frac{T \cdot \pi_w^*}{\exp(c_w/\sigma)}} = \frac{\sum_{w \in \mathcal{W}'} \pi_w^*}{\sum_{w \in \mathcal{W}'} \frac{\pi_w^*}{\exp(c_w/\sigma)}}. \quad (34)$$

Similarly, the average burst length of EconCast-C in anyput mode is computed by replacing c_w with γ_w in (34). Since $\gamma_w = 1$ always holds for $w \in \mathcal{W}'$, B_a simplifies to

$$B_a = \frac{\sum_{w \in \mathcal{W}'} \pi_w^*}{\sum_{w \in \mathcal{W}'} [\pi_w^* \cdot \exp(-1/\sigma)]} = \exp(1/\sigma). \quad (35)$$

This shows that the anyput average burst length is independent of the number of nodes, N , and only depends on σ . ■