Performance Evaluation of Energy-Constrained Broadcast (EconCast) in Wireless Networks

Tingjun Chen*, Javad Ghaderi*, Dan Rubenstein†, Gil Zussman*
*Electrical Engineering and †Computer Science, Columbia University, New York, NY 10027, USA
{tingjun, jghaderi, gil}@ee.columbia.edu, danr@cs.columbia.edu

***Abstract*—Wireless object tracking applications are gaining popularity and will soon utilize emerging ultra-low-power device-to-device communication. However, severe energy constraints require much more careful accounting of energy usage than what prior art provides. To address these challenges, in [1] we focused on maximizing broadcast throughput between a set of heterogeneous ultra-low-power devices and presented the Energy-Constrained Broadcast (EconCast) protocol. EconCast is a simple asynchronous distributed protocol that can operate in *groupput* and *anyput* modes to respectively maximize two alternative throughput measures. In this paper, we evaluate the throughput performance of EconCast numerically and via extensive simulations in various heterogeneous and homogeneous networks. We also evaluate the latency performance of EconCast and compare its throughput-delay tradeoffs when operating in groupput and anyput modes.**

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

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 in healthcare, smart building, assisted living, manufacturing, supply chain management, and intelligent transportation.

A fundamental challenge in networks of ultra-low-power nodes is to schedule the nodes' sleep, listen/receive, and transmit events without coordination, such that they communicate effectively while adhering to their strict power budgets. For example, energy harvesting tags need to rely on the power that can be harvested from sources such as indoor-light or kinetic energy, which provide between $0.01\,\mathrm{mW}$ and $0.1\,\mathrm{mW}$ [7]–[9] (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 [10] and ZigBee/802.15.4 [11] (usually at the order of $1-10\,\mathrm{mW}$). On the other hand, Bluetooth Low Energy and ZigBee are designed to support data rates (up to a few Mbps) that are higher than required by the applications our work envisages supporting (less than a few Kbps). Therefore, for such networks, the severe power constraints, the differing power consumption levels for listening, receiving, and transmitting, as well as the limited control bandwidth must all be considered.

In [1], we studied the problem of *maximizing broadcast throughput among energy-constrained nodes* and presented efficient methods to compute the maximum achievable throughput (i.e., oracle throughput). We also presented EconCast: **E**nergy-**con**strained Broad**Cast** – an asynchronous distributed protocol in which nodes transition between sleep, listen/receive, and transmit states, while maintaining a power budget, and can obtain throughput that approaches the maximum possible. *In this paper, we evaluate the throughput and latency performance of* EconCast *when operating in two modes: groupput mode and anyput mode*. These modes aim to maximize two alternative measures of broadcast throughput as defined below:

- **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]–[16]. 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.
- **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 [17]–[20], or decentralized coding [21].

We evaluate EconCast numerically and via extensive simulations for various heterogeneous and homogeneous networks. We show that the throughput achieved by EconCast approaches the oracle throughput in a limiting sense. We also evaluate the latency performance of EconCast and consider the throughput-delay tradeoffs of EconCast when operating in groupput and anyput modes.

The rest of the paper is organized as follows. After discussing related work in Section II, we present the problem formulation in Section III and describe the EconCast protocol in Section IV. We then evaluate EconCast numerically and via extensive simulations in Section V and conclude in Section VI.

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 (e.g., [12], [14]–[16], [22]–[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. Additionally, these protocols often require some explicit coordination (e.g., slotting [16], [24], or exchange of parameters [13]), which are not suitable for emerging ultra-low-power nodes. Moreover, many prior approaches achieve throughput levels which are much lower than optimal.

III. MODEL AND PROBLEM FORMULATION

To evaluate the *distributed* protocol EconCast, it is essential to compare its performance to the maximum achievable throughput (i.e., oracle throughput). In this section, we follow the model in [1] and describe the 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 $\mathcal{N}$.

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 values are 0, L_i (W), and X_i (W).[2] These power consumption levels are based on hardware characteristics.

Power budget: Each node i has a constant *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. 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. Also, 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 $j \neq i$.

B. Network Model and Throughput Definitions

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. We assume that: (i) the network is a *clique*,[3] (ii) there is only *one frequency channel* and *a single transmission rate* is used by all nodes in the transmit state, and (iii) nodes perform *perfect carrier sensing* prior to attempting transmission to check the availability of the medium, 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).

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 associated with a given network state $\mathbf{w}$ are defined below.

Definition 1: 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.,

$$\mathcal{T}_g = \lim_{T\to\infty} \frac{1}{T} \int_{t=0}^{T} \nu_{\mathbf{w}(t)} c_{\mathbf{w}(t)} \mathrm{d}t. \tag{1}$$

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

$$\mathcal{T}_a = \lim_{T\to\infty} \frac{1}{T} \int_{t=0}^{T} \nu_{\mathbf{w}(t)} \gamma_{\mathbf{w}(t)} \mathrm{d}t. \tag{2}$$

Definition 3: The *throughput associated with a given network state* $\mathbf{w} \in \mathcal{W}$, denoted by $\mathcal{T}_{\mathbf{w}}$, is defined as

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

C. Problem Formulation and Oracle Throughput

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

$$\pi_{\mathbf{z}} = \lim_{T\to\infty} \frac{1}{T} \int_{t=0}^{T} \mathbf{1}_{\{\mathbf{w}(t)=\mathbf{z}\}} \,\mathrm{d}t, \tag{4}$$

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

---

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

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

[3] We also investigated non-clique networks in [1].

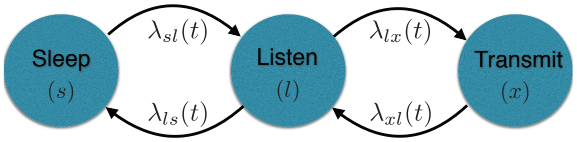

Fig. 1: The node's states and transition rates.

in sleep state is simply $(1-\alpha_i-\beta_i)$. In view of (1) – (4), (P1) is given by

$$\text{(P1)} \ \max_{\boldsymbol{\pi}} \ \sum\nolimits_{\mathbf{w}\in\mathcal{W}} \pi_{\mathbf{w}} \mathcal{T}_{\mathbf{w}} \tag{5}$$

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

$$\alpha_i = \sum\nolimits_{\mathbf{w}\in\mathcal{W}_i^l} \pi_{\mathbf{w}}, \ \beta_i = \sum\nolimits_{\mathbf{w}\in\mathcal{W}_i^x} \pi_{\mathbf{w}}, \tag{7}$$

$$\sum\nolimits_{\mathbf{w}\in\mathcal{W}} \pi_{\mathbf{w}} = 1, \ \pi_{\mathbf{w}} \ge 0, \ \forall \mathbf{w} \in \mathcal{W}, \tag{8}$$

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 grouput* and *oracle anyput* as $\mathcal{T}_g^*$ and $\mathcal{T}_a^*$. As mentioned above, in [1] we developed efficient methods to centrally compute the oracle throughput.

IV. THE ENERGY-CONSTRAINED BROADCAST PROTOCOL

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

As depicted in Fig. 1, a node can be in one of three states: sleep (s), listen (l), and transmit (x), and it must pass through the listen state to transition between sleep and transmit states. The time duration a node spends in a given state u before transitioning to state v is exponentially distributed with rate $\lambda_{uv}(t)$. Essentially, EconCast determines, in a distributed manner, how these transition rates are adjusted over time to achieve the best throughput performance.

Roughly speaking, each node adjusts its transition rates $\lambda_{uv}(t)$ based on limited information obtained in practice including its power consumption levels, L and X, and energy storage level, $b(t)$, a sensing of transmit activity of other nodes over the channel (CSMA-like carrier sensing), and a count of other active listeners (for groupput maximization), $c(t)$, or an indicator of whether there are any active listeners (for anyput maximization), $\gamma(t)$. In practice, $c(t)$ and $\gamma(t)$ may not be accurate, and we denote $\hat{c}(t)$ and $\hat{\gamma}(t)$ as their estimated values. We note that in EconCast, unlike in previous work such as Panda [13], Birthday [16], or Searchlight [15], 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.

To maximize groupput and anyput, EconCast can operate in *groupput mode* and *anyput mode*, respectively. EconCast takes input of three internal variables at any time t: (i) a multiplier $\eta(t)$ that is updated based on the energy availability of the node (i.e., as a function of its energy storage level $b(t)$), (ii) a carrier sensing indicator $A(t)$, which is 1 when the node does not sense any ongoing transmission, and is 0 otherwise, and (iii) a constant parameter $\sigma > 0$, which controls the achievable throughput. Using these variables, the transition rates of a node running EconCast (see Fig. 1) are described as follows, in which the two throughput modes only differ in $\lambda_{xl}(t)$. For groupput maximization,

$$\lambda_{sl}(t) = A(t) \cdot \exp[-\eta(t)L/\sigma], \tag{9a}$$

$$\lambda_{ls}(t) = A(t), \tag{9b}$$

$$\lambda_{lx}(t) = A(t) \cdot \exp[\eta(t)(L-X)/\sigma], \tag{9c}$$

$$\lambda_{xl}(t) = \exp[-\hat{c}(t)/\sigma]. \tag{9d}$$

For anyput maximization, $\hat{c}(t)$ is replaced with $\hat{\gamma}(t)$. We use $\mathcal{T}^\sigma$ ($\mathcal{T}_g^\sigma$ and $\mathcal{T}_a^\sigma$) to denote the analytical throughput (groupput and anyput) achieved by EconCast under a given value of σ.

V. NUMERICAL RESULTS

In this section, we evaluate the performance of EconCast when operating in groupput and anyput modes. For brevity, we ignore the subscripts of $\mathcal{T}^\sigma$ when describing results that are general for both groupput and anyput.

A. Setup

We consider clique networks and $\sigma \in \{0.1, 0.25, 0.5\}$. The nodes' power budgets and consumption levels correspond to energy harvesting budgets and ultra-low-power transceivers in [7], [8], [26]. 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 assume that the packet length is $1\,\mathrm{ms}$ and that nodes have accurate estimate of the number of listeners or the indicator of existence of active listeners, i.e., $\hat{c}(t) = c(t)$ or $\hat{\gamma}(t) = \gamma(t)$.

Our results show that the simulated throughput perfectly matches the analytical throughput $\mathcal{T}^\sigma$ for $\sigma \in \{0.25, 0.5\}$. For $\sigma = 0.1$, the simulated throughput does not converge to $\mathcal{T}^\sigma$ within reasonable time due to the bursty nature of EconCast, as will be described in Section V-D. Therefore, we evaluate the throughput performance of EconCast by comparing $\mathcal{T}^\sigma$ 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.

B. Throughput in Heterogeneous Networks

One strength of EconCast is its ability to deal with *heterogeneous* networks. Fig. 2 shows the groupput and anyput

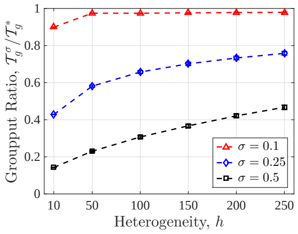

(a) (b)

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.

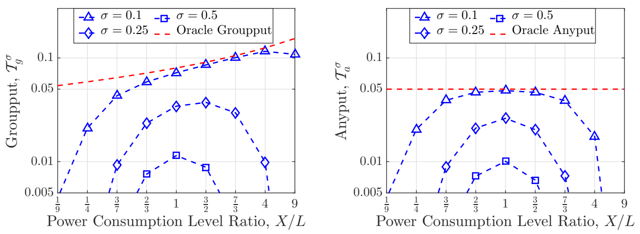

(a) (b)

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.

achieved by EconCast normalized to the corresponding oracle groupput and anyput, for heterogeneous networks with $N = 5$ and $\sigma \in \{0.1, 0.25, 0.5\}$. Intuitively, higher values of the throughput ratio $\mathcal{T}^\sigma/\mathcal{T}^*$ indicate better performance of EconCast. To explicitly describe the results in a controlled manner, we represent the *network heterogeneity* by one single parameter h. Along the x-axis, we vary h 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), and are with mean values of $500\,\mu\text{W}$, (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 power budget ρ_i varies from $100/h$ to h (μW) and has median of $10\,\mu\text{W}$, but its mean increases as h increases.

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 $\sigma \to 0$. 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 ratio and anyput ratio achieved by EconCast have similar performance except for homogeneous networks ($h = 10$), in which the anyput ratio is slightly higher than the groupput ratio. This is because that in homogeneous networks, nodes operating in a fully distributed manner are easier to collect the information of $\gamma(t)$, and therefore, $\gamma(t)$ is more important for improving the throughput than $c(t)$.

C. Throughput in Homogeneous Networks

We now evaluate the groupput and anyput performance in *homogeneous* networks with $N = 5$, $\rho = 10\,\mu\text{W}$, $L + X = 1\,\text{mW}$, and $\sigma \in \{0.1, 0.25, 0.5\}$. Figs. 3(a) and 3(b) present, respectively, the groupput and anyput achieved by EconCast as a function of the power consumption ratio X/L. In addition, the top dashed lines represent the oracle groupput and anyput. Fig. 3 shows that with fixed sum of power consumption levels (i.e., $L + X = 1\,\text{mW}$), the oracle groupput $\mathcal{T}_g^*$ increases as X/L increases, but the oracle anyput $\mathcal{T}_a^*$ remains constant. This is consistent with our groupput and anyput analysis in [1]:

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

For a given value of X/L, $\mathcal{T}^\sigma$ approaches $\mathcal{T}^*$ as σ decreases. 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 (i.e., $L \approx X$), which is realistic for current commercial low-power radios that have symmetric power consumption levels in listen and transmit states.

It is observed that the throughput performance of EconCast degrades with extreme values of X/L. This is because (i) with small X/L values, nodes enter transmit state infrequently, since listen is expensive and they must pass the listen state to enter the transmit state, and (ii) with large X/L values, nodes waste their energy to transmit even when there is no other nodes listening (e.g., $c(t) = 0$). In particular, anyput degrades with large X/L values since anyput depends on the *existence* of listening nodes when some node is transmitting. Therefore, when listening is expensive, the fact that multiple nodes listen simultaneously does not improve anyput. However, we believe that any distributed protocol will suffer from such performance degradation since nodes in a fully distributed setting *do not* have any information about the properties of other nodes.

D. Latency

The results until now suggest allowing $\sigma \to 0$. While reducing σ improves throughput, it considerably increases both the *burstiness* and the communication *latency*. In general, increased burstiness means that the long term throughput can be achieved with given power budgets but the variance is more significant during short term intervals (for details, see [1]). Here, we focus on the latency of EconCast, 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 measure the latency compared to the unit packet length of $1\,\text{ms}$ as described in SectionV-A. Figs. 4(a) and 4(b) present the CDF latency of EconCast when operating in groupput and anyput modes obtained via simulations, and indicate both the average and the 99^{th} percentile latency values. The homogeneous networks considered

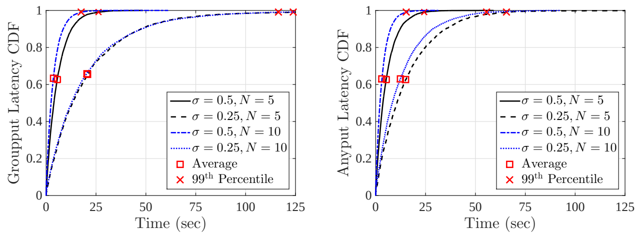

(a) (b)

Fig. 4: The CDF, mean, and 99^{th} 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 $L = X = 500\,\mu\text{W}$.

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. 4 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. 4 also illustrates that a larger value of N results in lower latency, since it is more likely to receive when more nodes exist. Comparing Fig. 4(a) with Fig. 4(b), it is observed that EconCast operating in anyput mode has slightly lower average latency than when operating in groupput mode. However, with a smaller value σ (i.e., $\sigma = 0.25$), the 99^{th} percentile latency of EconCast when operating in anyput mode is significantly lower than that in groupput mode. This is because that the burst length of EconCast (which heavily depends on σ) 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.

VI. CONCLUSION

In [1], we presented and analyzed EconCast, an asynchronous distributed protocol designed to maximize the broadcast throughput among a set of heterogeneous energy-constrained nodes. In this paper, we evaluated the throughput and latency performance of EconCast when operating in groupput and anyput modes. First, we numerically showed that the throughput (i.e., groupput and anyput) achieved by EconCast approaches the oracle throughput as $\sigma \to 0$, in various heterogeneous and homogeneous networks. We also studied the sensitivity of the achievable throughput to the network heterogeneity. Then, via simulations, we evaluated both the throughput achieved by EconCast and its latency performance. Finally, we considered and compared the throughput-delay tradeoffs of EconCast when operating in both groupput and anyput modes.

ACKNOWLEDGEMENTS

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

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, Oct. 2010.

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

[4] R. Margolies, M. Gorlatova, J. Sarik, G. Stanje, J. Zhu, P. Miller, M. Szczodrak, 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, pp. 62:1–62:27, Nov. 2015.

[5] 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.

[6] V. Liu, A. Parks, V. Talla, S. Gollakota, D. Wetherall, and J. R. Smith, “Ambient backscatter: wireless communication out of thin air,” in *Proc. ACM SIGCOMM’13*, 2013.

[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. Kymissis, 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, Aug. 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] 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.

[14] 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.

[15] M. Bakht, M. Trower, and R. H. Kravets, “Searchlight: Won’t you be my neighbor?” in *Proc. ACM MobiCom’12*, 2012.

[16] 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.

[17] J. Haupt, W. U. Bajwa, M. Rabbat, and R. Nowak, “Compressed sensing for networked data,” *IEEE Sigal Process Mag.*, vol. 25, no. 2, pp. 92–101, 2008.

[18] 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.

[19] T. Srisooksai, 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.

[20] M. Rabbat, J. Haupt, A. Singh, and R. Nowak, “Decentralized compression and predistribution via randomized gossiping,” in *Proc. ACM/IEEE IPSN’06*, 2006.

[21] 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.

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

[23] 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.

[24] A. Purohit, B. Priyantha, and J. Liu, “WiFlock: Collaborative group discovery and maintenance in mobile sensor networks,” in *Proc. ACM/IEEE IPSN’11*, 2011.

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