

EMIT: An Efficient MAC Paradigm for the Internet of Things

Arjun Bakshi^{*†}, Lu Chen^{*†}, Kannan Srinivasan[†], C. Emre Koksal[‡], Atilla Eryilmaz[‡]

[†]Department of Computer Science and Engineering, The Ohio State University, Columbus 43210, OH

[‡]Department of Electrical and Computer Engineering, The Ohio State University, Columbus 43210, OH

^{*}Co-primary authors

Abstract—The future Internet of Things (IoT) networks are expected to be composed of a large population of low-cost devices communicating dynamically with access points or neighboring devices to communicate small bundles of delay-sensitive data. To support the high-intensity and short-lived demands of these emerging networks, we propose an Efficient MAC paradigm for IoT (EMIT). Our paradigm bypasses the high overhead and coordination costs of existing MAC solutions by employing an interference-averaging strategy that allow users to share their resources simultaneously. In contrast to the predominant interference-suppressing approaches, EMIT exploits the dense and dynamic nature of IoT networks to reduce the spatio-temporal variability of interference to achieve low-delay and high-reliability in service.

This paper introduces foundational ideas of EMIT by characterizing the global interference statistics in terms of single-device operation and develops power-rate allocation strategies to guarantee low-delay high-reliability performance. A significant portion of our work is aimed at validating these theoretical principles in experimental testbeds, where we compare the performance of EMIT to a CSMA-based MAC protocol. Our comparisons confirm the beneficial characteristics of EMIT, and reveal significant gains over CSMA strategies in the case of IoT traffic.

I. INTRODUCTION

By 2020, there will be an estimated 50 billion devices that will connect to the Internet. These Internet of Things (IoT) devices will cater to applications like smart homes, body/health monitoring, environmental monitoring, condition-based maintenance, among many others. Despite the significant attention, the IoT concept is fairly young. IEEE Standards Association (IEEE-SA) has recently created a working group to outline the architecture needed to support IoT. IoT’s architecture is quite open. This is an opportune time to explore a few possibilities to influence the standardization efforts.

In many applications targeted for IoT, it’s expected that the per-station traffic will be sparse/intermittent in time whereas stations will be dense in space. Take, for example, a smart home application. An intrusion monitoring system will have many sensors deployed throughout the house, along with sensors in home appliances monitoring generating other data. Thus, these sensors will be deployed densely, while each sensor will intermittently report its current status.

These characteristics contrast sharply with their wireless local area network (WLAN) counterparts in which intense traffic is generated by relatively sparsely positioned

stations. Since typical transmissions for such a scenario will be long and the number of stations is low, designing multiple access, based on careful coordination to decide on what resources to assign to which users highly improves system performance. Accordingly, it is worth spending time and energy to orthogonalize transmissions over time, frequency, and/or code to achieve high performance. Thus, many existing WLAN multiple access schemes (CDMA, OFDMA, CSMA) are designed towards this end.

However, for dense IoT devices with intermittent traffic demands, it becomes proportionally costly to coordinate transmissions for each new bundle of data service. This is not justified, especially given that per-device traffic is sparse in time. Furthermore, under CSMA, nodes sense the channel and use backoffs to avoid collisions (or interference), which typically leads to unreasonable amounts of delay in scenarios under high node densities per access point (e.g., as we all experience in a crowded Starbucks to access the web.). Such long delays will affect applications that require timely delivery of small, but intermittent data.

These observations lead to the question: *Why endure a high-overhead and large-delay MAC protocol in IoT networks if we are going to send a few packets that arrive intermittently to each device?* In answer to this question, we propose an interference-embracing paradigm in which we allow many users to share the resources simultaneously. While this bypasses the heavy costs of orthogonalization, the important question is what we do about the resulting co-channel interference that can be extremely harmful for network performance? In response to this challenge, our design exploits the combination of spatial density of devices with temporal sparsity of traffic to create *statistical certainties*. The aim here is not to reduce interference to zero, rather to push the variability of interference to zero with increasing network size.

To achieve this, whenever IoT devices have traffic, they transmit immediately but at carefully chosen transmission durations and power levels. These choices¹ of duration and power, combined with dense deployment and intermittent nature traffic, results in a sum interference that is more stable across space and time, a phenomenon we refer to as **interference averaging**. This “predictability” of interference allows each transmitter to achieve high performance

¹Section III outlines how transmission duration and power are chosen and their relationship to statistical certainty.

at the physical layer, since it can choose a *robust code that does not have to be conservative to achieve low drop-rates*. Consequently, we observe higher *achieved area spectral efficiency* with EMIT, compared to CSMA. Furthermore, the complexity of higher layers reduces as the need for centralized or distributed coordination is mostly avoided. Note that, interference averaging enables the system to *handle variations due to user dynamics and data arrivals, directly at the physical layer*, eliminating the need for solutions based on careful scheduling of user transmissions.

Overall, this paper makes the following contributions;

- We lay down the theoretical foundations of a new network architecture geared towards efficient IoT network operation in Section III. This results in an interference averaging paradigm with high-reliability and low-delay service guarantees for IoT services with intermittent demands.
- We present a completely asynchronous and distributed solution for data communication across IoT, encompassing MAC and physical layers that consists of a set of simple coding, power, and rate control strategies in Section III.
- We present a simplified implementation of EMIT and multiple extensive experimental results from different testbeds in Section IV.

II. MAIN IDEAS AND INSIGHTS

Here, we discuss the main insights and guiding principles of EMIT which motivate our design choices presented in the next section, where we formalize these insights.

• *What are the key characteristics of IoT networks?*

IoT networks are expected to be composed of a large population of low-cost devices communicating dynamically with access points or neighboring devices to communicate small bundles of typically delay-sensitive data.

Accordingly, IoT demands generated by densely packed mobile devices will come with an extraordinary collective activity, formed by the intermittent, delay-sensitive, and short-lived communications *per device* within a small area. These characteristics call for the development of low-complexity operation principles that enable low delay, low overhead, and reliable communication of IoT devices in densely packed environments.

• *What is wrong with the use of existing wireless technologies for IoT services?*

The traditional approach to MAC is to orthogonalize users in resource space (frequency, time, space) in order to reduce the interference at an active receiver during a transmission. This principle is justified for the service of consistent demands of a small-to-moderate number of users in a wireless service area where the overhead introduced by backoffs, or coordinations is negligible compared to the payload being transmitted. In contrast, in IoT networks, due to the increased device density, the cost of coordination and delay from backoffs will become much greater than the cost of transmitting the payload[1].

• *What are the core principles behind EMIT?*

In our envisioned scheme, EMIT, each device with data in its queue starts its transmission at a sufficiently low power level and a carefully chosen data rate. Due to the low transmission power, data is encoded across a long duration. Multiple such simultaneous transmissions from a large number of devices at low power and low rate, combined with intermittency and sparsity of traffic leads to a law-of-large-numbers-like averaging behavior for the observed interference power over entire time and space.

Figs. 1(a) and 1(b) depict a typical snapshot of observed interference at every point in a network under CSMA and EMIT ². It can be seen that compared to CSMA, in addition to lower interference variability, EMIT also reduces the mean interference, averaged across space and time. The reason for this reduction is the convexity of the power-decay law for wireless propagation (in far field): a random point in space under CSMA will observe a very high interference if it is located close to a transmitter, or a small interference otherwise. In contrast, with EMIT, since power levels are low and transmitters are dense, the interference at a random point is less sensitive to the position. Using Jensen's inequality, higher variability of the distance from the active stations leads to a higher average interference.

The reduction in variability of interference due to *interference averaging* is vital in achieving high delay-limited rates for the short-lived IoT services. Note that the gain of this scheme is not merely due to simultaneous transmissions at a lower rate; instead, *most of the gain is associated with averaging of user/traffic dynamics at the physical layer*.

The notion of interference averaging is also exploited in spread-spectrum based systems, such as Direct Sequence Code Division Multiple Access (DS-CDMA). DS-CDMA spreads the symbols at the physical layer via pseudo-noise sequences to achieve statistical orthogonality across users, without considering user dynamics or traffic while interference averaging with EMIT is based on **user dynamics**. It exploits the intermittency of traffic in a way that, the averaging is observed at a time scale that averages out user-dynamics. It targets a dense network setting with sparse per node traffic to achieve averaging at a much longer time scale, compared to spread spectrum. Thus, EMIT eliminates the use of standard MAC protocols, unlike spread-spectrum based MAC schemes.

A key property of EMIT's design (elaborated in Section III) is the fact that if all device dynamics are short term, then the global interference will also be short term. This characteristic is the main enabler in shaping the global interference through individual IoT-device dynamics, and enables decentralized and efficient network-wide design. Consequently, the ultimate objective to achieve with EMIT is the behavior of reduced interference mean and variability

²The propagation model for this simulation uses a path loss component for power decay, combined with a fast fading component, independent and identically distributed across space. For a fair comparison, the total energy consumed by each device per communication/transmission is fixed, so shorter transmissions imply a proportionally higher transmission power.

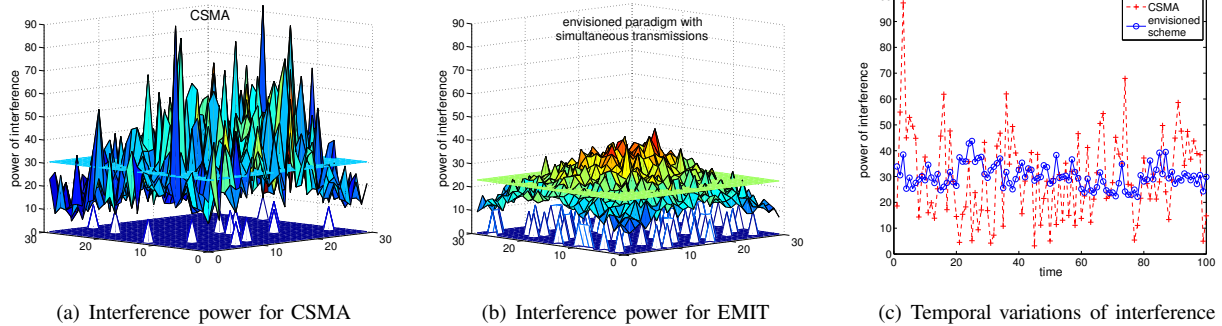


Fig. 1: Spatio-temporal behavior of the observed interference power. In (a) and (b), the active users are shown as triangles. In CSMA, while the active number of users at a time is low, the interference pattern across the space is more varied. In EMIT, the number of active devices is large but each device emits low power, leading to the **interference averaging** phenomenon. Similarly, as shown in (c), the temporal variability of the interference also reduces under EMIT.

both in time and in space. This leads to an increase in the area spectral efficiency, measured in bits/sec/Hz/m², which is an indicator for the achievable cumulative rate over a network. Next we show the significance of this principle for high-satisfaction IoT services.

• **Why are the new principles better-suited for IoT services than existing technologies?**

The interference-averaging principle has several desirable network-level performance implications for IoT services. First, the parallel transmission approach can provide better short-term guarantees for all IoT devices than the interference-avoiding protocols, like CSMA, that incur variable and large delays for service. Second, low interference-variability enables a device to choose a robust data rate, without being conservative to reduce packet drops due to interference spikes seen in interference avoidance protocols. Consequently, one can design modulation and code selection less conservatively for the *entire activity period*. This increases the area spectral efficiency in such a way that the message delay is reduced significantly, without sacrificing throughput. Third, this design enables decentralized operation of large-scale IoT applications with low coordination and overhead costs between devices as it allows many asynchronous transmissions within interference regions.

III. THEORY AND DESIGN OF EMIT

This section presents a sequence of analytical findings in order to support the design choices made based on the insights provided in the previous section. Here, the aim is to design a low-overhead, low-delay, and fully decentralized channel access IoT-MAC strategy for EMIT.

A. Interference Characteristics under IoT Dynamics

Since the design goal is to develop a distributed, lightweight, low-delay MAC, the effect of a single node's transmission on the global ensemble interference is studied when every node transmits as soon as a packet arrives in its transmission queue. The following analysis presents the spatio-temporal correlation of global interference. It shows that the *global interference correlation is heavily dictated by the channel access dynamics of a single node*.

a) Single Device MAC shapes Global Interference Statistics: Consider a 2-D plane that is densely filled with IoT devices with intermittent, short-lived service demands. To capture such dynamics in an analytical model, suppose IoT demands arise according to a spatio-temporal Poisson process³ over the entire 2D space with intensity λ demands/s/m². Each device starts transmitting as soon as a demand arrives, and remains actively transmitting at a constant power P for a random duration of time T seconds with distribution $F_T(t)$. The signal power observed by its intended receiver, originating from a node transmitting at power P located at a distance r is given by $H(r)P$, where $H(r)$ captures a random power gain (from the wireless channel) at a distance of r . It is assumed that there is no collision avoidance or scheduling, and the users operate completely asynchronously.

Let the interference power observed at the origin be denoted as $Z(t)$, which is statistically identical to the interference power observed at any other position in space. This results in the following second-order characterization of $Z(t)$ (see the technical report [2] for the derivation): the auto-covariance function of $Z(t)$ is

$$K_Z(t - \tau) = \underbrace{\left(\int_0^\infty 2\pi r P^2 (\mathbb{E}[H(r)])^2 dr \right)}_{\text{Governed by Spatial Statistics}} \cdot \underbrace{\left(\int_{|t-\tau|}^\infty \lambda [1 - F_T(s)] ds \right)}_{\text{Governed by Single-User Activity}}. \quad (1)$$

Interestingly, the expression reveals a natural decoupling between the impact of spatial elements and temporal elements of the setup. The spatial component concerns only the channel gain distribution and the transmit power level, while the temporal component captures the effect of user activity on the correlation structure. One may think that the collective effect of infinitely many interferers may lead to a temporal correlation structure that extends beyond the characteristics of a single user activity. However, Eq. (1)

³The number of IoT devices that become active within any unit area is a temporal Poisson process and the locations of the users is a 2D spatial Poisson process.

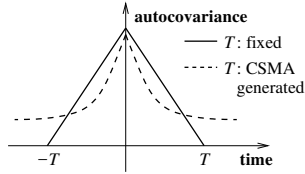


Fig. 2: Interference autocovariance functions under fixed vs. CSMA-generated transmission durations.

reveals that **a single node’s dynamics dictate the correlation of interference**. Indeed, the power spectral density of the cumulative interference has the same “shape” as that of the interference caused by a single user (cf. Fig. 2). In particular, the power spectral densities of the cumulative and the single user interference have the same bandwidth.

Figure 2 shows the auto-covariance of global interference for the case when a single node uses a fixed transmission duration to send its data (solid line). It shows that the global interference is also correlated for the same duration. The same figure shows global interference for the CSMA case (dashed line). Since CSMA randomly backs off before it transmits, the correlation duration is significantly higher than the fixed duration case. This is happening because a transmission under CSMA delays the transmissions in its neighborhood thus, creating longer-term correlation. To see this more clearly, one can relate the interference to waiting times in queueing systems: our system in which all users transmit upon arrival corresponds to an $M/G/\infty$ queue, whereas a scheduling system could be related to an $M/G/k$ queue. In the former system, the waiting time (corresponding to the interference correlation time in this context) is merely the service time of a single user. Whereas in the latter system, the nodes experience additional queuing delay as they wait for a server to become available. Consequently, the interference under collision avoiding scheduling involves a combination of queuing delay and the service time. \square

B. Physical layer and Decentralized MAC Design

As noted in Section III-A, there is a strong tie between a single node’s transmission characteristics and the global interference. Specifically, when the node uses a deterministic duration to transmit without backing off, then the global interference is correlated for that duration. This shows that *EMIT can average out the affect of global IoT activity within the duration of a single activity period*. This is achievable, for example, by treating the short-lived activity period of a device as a single resource block. Second, the knowledge of temporal interference statistics (as derived in (1)) enables the design of new rate/power allocation and coding strategies to maximize the service rates within a typical activity period.

In this effort, it is critical to note that IoT applications require communication over a finite connection duration. This could be due to many reasons such as limited energy, validity of the current data, time-sensitivity of the data (alarm, video/audio stream), etc. Thus, we are in the delay-limited setting. Our goal is not to maximize the total throughput for every node. Note that CSMA,

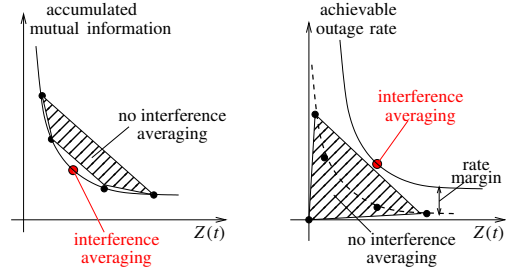


Fig. 3: Impact of interference variability on performance: for the same mean, interference variability improves the accumulated mutual information (left), while it can significantly hurt the achieved outage rate (right).

even after significant fine-tuning, provides interference-free opportunities for different devices at random times. During such durations, the SINR for that node is likely very high, allowing transmissions at very high data rates.

However, EMIT is interference-embracing and nodes send data as they arrive. Therefore, the SINR will be lower than it is in the CSMA case. However, the timely transmissions at a data rate commensurate with the observed SINR will minimize outages⁴. In order to reduce the probability of an outage, we need to put aside some **rate margin** and choose the data rate smaller than the expected accumulated mutual information. This margin has a critical impact on the performance. With this margin, the advantage of interference variability on the expected mutual information is negated by the occurrence of outages it causes. As illustrated on the right plot in Fig. 3, interference averaging starts to become favorable, as the margin size exceeds a certain threshold.

a) IoT-MAC Design for Outage Rate Maximization:

In EMIT, the transmission duration, T , is deterministic and each IoT device with a new service demand starts its transmission without any delay. In particular, when a new demand arrives, say at time t , the device senses the interference level $Z(t) = z(t)$ at the time and selects a power $P_{z(t)}$ and a rate $R_{z(t)}$ level, to be held over the transmission duration T . We assume that each device has a constraint on the average amount, E , of energy it consumes during any given transmission⁵. Since the transmission will experience an outage if the accumulated mutual information rate is lower than the transmission rate, the objective of the allocation becomes that of an **outage-rate maximization**:

$$\begin{aligned} \max_{P_{Z(t)}, R_{Z(t)} \geq 0} & \mathbb{E} \left[R_{Z(t)} \mathbb{I} \left(\int_t^{t+T} \right. \right. \\ & \left. \left. W \log \left(1 + \frac{P_{Z(t)}}{N_0 W + Z(\tau)} \right) d\tau \geq R_{Z(t)} \right) \right] \quad (2) \\ \text{s. t.} & \quad T \cdot \mathbb{E}[P_{Z(t)}] \leq E, \quad (3) \end{aligned}$$

⁴By outage, we mean that the data has a deadline before which it needs to be sent. Beyond that deadline, that data is not useful.

⁵In the IoT context with short-lived connections, an energy constraint over the connection periods is more relevant, as opposed to the more-traditional average power constraint, since the interconnection times are much larger than the delay constraints of the applications generating the connections.

where $\mathbb{I}(\cdot)$ is the indicator function, (2) is the expected amount of information that does not experience an outage, over a transmission, and (3) is the average energy constraint over different transmissions.

Figs. 4(a) and 4(b) show the cumulative information rate after solving the above equation [3], [4], [5], [6], [7]. The following observations are in order;

Take-Aways: (i) We observe that as T increases, EMIT allocates increasingly higher portion of its energy and rate to lower interference levels. This behavior is de-mystified once we recall the impact of T on the shape of the autocovariance of interference exposed in Section III-A: as T increases so does the correlation time of interference. This implies, in turn, that it is more efficient to allocate the limited energy for transmissions when the interference level is lower, since its level is expected to remain low for a significant portion of the transmission duration. (ii) For any given T , there is a point in the observed interference level, beyond which the power allocation goes to 0. This cutoff point decreases as T increases due to *interference averaging* phenomenon we illustrated in Fig. 1.

In Fig. 4(c), we illustrate the cumulative amount of transmitted information per transmission that does not experience an outage (along with the accumulated mutual information that does not account for outages) under the above selection of $(P_{z(t)}, R_{z(t)})$. Note that, average SNR decreases with T (from 27 dB for $T = 1$ to 7 dB for $T = 100$ for the conditions with typical power levels), since the constraint (3) is not on power, but on energy per transmission. Despite the decrease in SNR (and hence the spectral efficiency), the accumulated mutual information is monotonically increasing with T as shown in Fig. 4(c) due to the interference averaging nature of EMIT. \square

C. EMIT Algorithm Description

The above observations, based on the model and the line of analysis provided, have important consequences in the design of EMIT. It shows that, since all user dynamics are short-term, the global interference will also be short-term. With EMIT, when each user encodes the associated data as a single block and transmits it directly at an appropriately chosen power and rate pair, then a longer transmission duration is preferable. Note that, this expanded duration is not merely due to lowered rate associated with lower SINR. Instead, the main objective of this slow-down is to **average out user dynamics at a single encoding block**, so that such dynamics are handled at the physical layer, bypassing MAC. It is also observed that, one should choose T as large as the applications QoS requirements allow. This statement is somewhat surprising, since the SINR decreases as T^{-1} . In Section IV, we show that the achieved area spectral efficiency⁶ measured in bits/sec/m² that EMIT achieves is superior to that achieved by CSMA; and EMIT

⁶Area spectral efficiency gives a direct measure for the amount of bits each user can transmit during the activity period. To understand this, simply divide the area spectral efficiency A bits/sec/Hz/m² by the constant λ users/sec/m² to obtain A/λ bits/Hz/user, the achieved rate per user over unit bandwidth.

Algorithm 1: EMIT Algorithm at each IoT device

```

1 Input:
2 Encoding block length  $T$ 

3 Operation state:
4 while Message queue non-empty at time  $t$ , do
5   observe interference power  $z(t)$ 
6   set power  $P_{z(t)}$  and rate  $R_{z(t)}$ , using the solution of
   (2,3).
7   transmit head-of-line message at power  $P_{z(t)}$ , encoded
   at rate  $R_{z(t)}$  over the current block.
8 endwhile

```

uses this increase in such a way that the delay is reduced significantly, without reducing the throughput.

The pseudo-code for our algorithm is given in Algorithm 1. As can be seen, the algorithm is fully decentralized and demands no synchronization among ongoing transmissions. The only synchronization requirement is between each transmitter-receiver pair (frame and symbol) to decode transmitted messages. To that end, frame detection and synchronization is achieved by using a suitably long preamble. Note that due to simultaneous transmissions directed to the same receiver, there are ways to make sure these preamble sequences work in the low SINR regimes [8].

IV. EXPERIMENTAL EVALUATION

This section presents experiments and results generated from two testbeds to validate the ideas presented in the previous sections, and to test the coexistence of EMIT with existing predominant MAC technologies.

A. Testbeds and Experiments

Experiments are conducted on two testbeds: a USRP testbed at Orbit Lab at Rutgers [9], and a local testbed constructed using TelosB motes. TinyOS and nesC are used to program the motes, and GNU Radio for the USRPs.

1) *TelosB testbed:* The TelosB testbed is setup in a classroom using 36 motes. Motes used for creating interference are placed in groups of 3 and 4 as a grid around a receiver mote and USRP. This setup covers an area of approximately 6 meters by 6 meters. The receiving mote is used to measure the packet reception rate, while the USRP is used to record RSSI and signal.

2) *Orbit Lab USRP testbed:* The TelosB testbed is a relatively dense testbed with approximately 1 mote/m². The Grid testbed at Orbit Lab is much sparser and can provide insights into how EMIT performs when network density is reduced. The testbed is a grid of size 20 meters by 20 meters, with USRPs spread over the grid. The experiments use 23 Ettus USRP N210 with SBX daughterboards.

B. Experiment variables

In addition to the varying size and density of the network, the following parameters are also varied:

Packet Arrival Rate, λ : A Poisson random variable used to control the message arrival rate at every node. It is set to 2, 4, 8, and 16 (packets/sec) in the experiments.

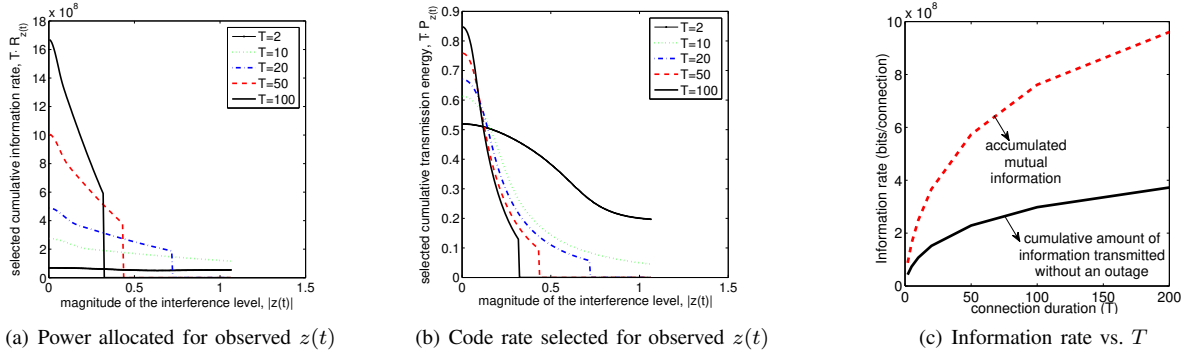


Fig. 4: In (a) and (b), we show optimal power and rate allocations respectively, as a function of the interference signal magnitude, $|Z(t)|$, for different transmission durations T . In (c), accumulated mutual information and achievable cumulative rate are given.

Packet transmission duration (T_{tx}) and transmission power (P_{tx}): A simplified version of EMIT is implemented where the transmission power P is selected as a constant independent of the interference level $Z(t)$. This variation alleviates the need for interference measurement and feedback, and serves as a lower bound on the performance of full EMIT. In addition, the transmission and duration are varied in tandem so that the total energy consumed during packet transmission remains the same across different settings. These settings are referred to as equal energy settings and are represented with a pair of the form (P_{tx}, T_{tx}) .

The equal energy settings chosen for the TelosB motes are (0 dBm, T), (-3 dBm, 2T), (-7 dBm, 5T). Lower power setting could not be tested for the motes due to hardware restrictions. Since the USRPs are not calibrated, the transmission power is not reported in dBm, and using Tx gain instead. The equal energy settings chosen for the USRPs are (12, T), (9, 2T), (6, 4T), (3, 8T) and (0, 16T).

C. Transmission Scenario

In the experiments, all the transmitters are trying to communicate with one receiver. This can be thought of as a gateway node that connects a group of IoT like devices to a larger network. Each device tries to transmit a fixed amount of data, $bytes_{data}$, to the receiver by transmitting for a fixed amount of time. That is, the transmission rate is determined by $bytes_{data}/T$. However the coding rate $R_{z(t)}$ is varied based on the SINR observed, $z(t)$, by the transmitter as dictated by the EMIT protocol. The performance of EMIT is compared to CSMA⁷ under this setup, however this restriction is later dropped to take a closer look at where the strengths and weaknesses of EMIT lie.

D. Testing Interference Averaging of EMIT

Section II argued that the net interference observed by any node in a network becomes more stable under the EMIT paradigm as every node chooses a larger time duration, T_{tx} , and a lower transmission power P_{tx} such that the total energy $E = P_{tx} \cdot T_{tx}$ is kept the same. A key question to ask is: **Is interference averaging observable in practice?**

⁷We use the same CSMA scheme as IEEE 802.15.4 in both testbeds.

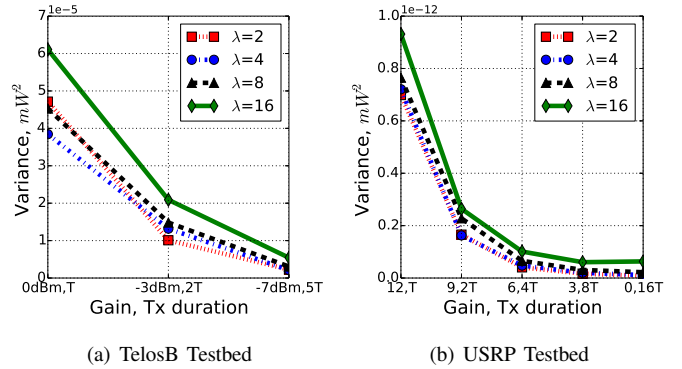


Fig. 5: Observed interference's variance. The variance decreases as the transmission duration increases.

To this end, the variance of interference for all equal energy settings and different λ is computed.

1) *Variance of Interference:* The variance of interference for the two testbeds seen at a receiver is shown as a function of λ and equal energy settings in Figures 5(a) and 5(b). It is clear that as the transmission duration increases, the variance of interference indeed decreases as predicted. This trend is seen for all values of λ considered.

One main advantage of interference averaging is that a transmitting node can choose a reasonable data rate that is commensurate with the observed interference. In the absence of interference averaging, interference may increase during the packet transmission itself and therefore, a transmitter has to choose a rate that is conservative.

Data rate is not just a function of the observed interference, it is rather a function of the SINR. Note that SINR's variance is also affected by the variance of channel during transmissions. However, coherence time – the time for which the channel remains stationary or correlated – tends to be stable in the order of hundreds of milliseconds. It can be inferred from the interference variance results in Figures 5(a) and 5(b) that for a fixed power signal, the SINR will also follow the same trend.

E. Comparing Delays of EMIT and CSMA

A major claim in previous sections is that the delays in transmitting packets will be reduced in the EMIT paradigm compared with CSMA, which will spend more time per-

forming channel sensing and back-off. Experiments to examine the delays of EMIT and CSMA, considering different packet arrival rates, and packet transmission durations are conducted. On both testbeds, CSMA is run with the shortest transmission duration T_{tx} , while EMIT is run for different equal energy settings.

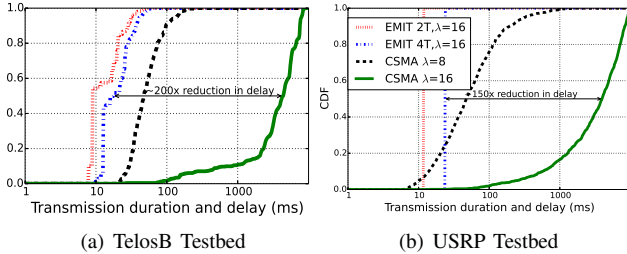


Fig. 6: CSMA vs EMIT: Delay Performance.

Figure 6(b) shows that the delay time for EMIT is almost constant even for the highest λ setting, while that of CSMA grows as λ increases. This is because as λ increases, CSMA needs to back-off more often. The only case where packets in EMIT will have to wait in the queue is if the packet inter-arrival time is comparable to the time EMIT takes to transmit a packet, which is the reason behind the variable delay seen in the delay curves of EMIT in Figure 6(a). Even for the latter scenario, the delay under EMIT is much lower than that of CSMA. Also note that the delay experienced by EMIT is not dependent on the number of nodes in the network, as would be the case for CSMA.

Clearly, EMIT’s delay performance is significantly better than that of CSMA. However, **is the gain in delay coming from a loss in throughput?** The following section examines this question.

F. Comparing Throughput of EMIT and CSMA

The reduction in delay time for EMIT over CSMA comes from immediately transmitting whatever rate the channel is expected to support (based on calculation in Eq.2). Although this will cause EMIT to transmit at a much lower rate than CSMA, this disadvantage is offset by two factors. Firstly, since all devices will follow the same procedure for transmission, interference averaging will cause the observed interference to be more steady. This means it is more likely that the rate supported by the channel at the beginning of the transmission is likely to stay the same for the entire transmission duration. Secondly, although the SINR will be lower than the SINR observed by CSMA, the throughput depends logarithmically on the SINR while linearly on the transmission duration. Therefore the loss from lower SINR is offset, to some extent, by the increase in transmission duration.

1) *Throughput Comparison under Fixed Data per Message:* The scenario examined here is, given a packet with a payload of size $bytes_{data}$, how often is EMIT able to transmit the packet using an equal energy setting without suffering from outages? This experiment is aimed at assessing the utility of EMIT for an IoT-like setup, where the payload size is fixed for all or most transmissions, and

the message arrival times are periodic. This is in contrast to settings where the amount of data to the transmitted changes on a per packet basis.

To do this, the cumulative throughput of each transmission is calculated based on the Shannon rate and the SINR observed during the transmission duration. If this is greater than the size of the packet to be transmitted, then transmission is successful, otherwise it is considered to be a failed transmission. This is repeated for all values of λ .

We tried to send 1000, 5000 and 10000 bits/message for both schemes, and the results showed that EMIT, under all equal energy settings and λ , is always able match the packet success rate of CSMA. Next we examine the unbounded throughput of EMIT which explains this result.

2) *Maximum Achieved Throughput Comparison:* This section examines the throughput of a typical link under CSMA and EMIT using Shannon’s capacity equation when transmissions are not restricted to a fixed packet size.

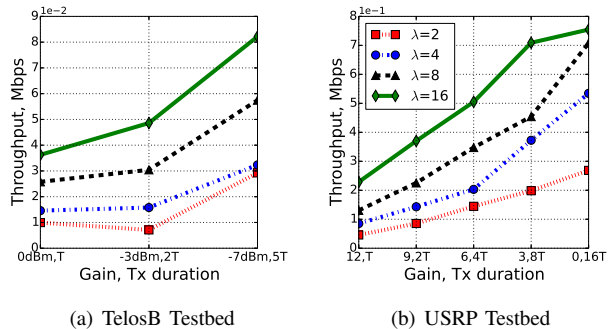


Fig. 7: Throughput of EMIT

Figures 7(a) and 7(b) show that as the transmission duration increases (and transmission power decreases) under EMIT, the throughput of a device increases. This shows that EMIT can achieve throughput demands despite operating at low SINR.

A direct comparison between the throughput of CSMA and EMIT both operating at the highest power setting for all values of λ is shown in Figure 8. It shows that CSMA comes very close to transmitting at the channel capacity. Its loss comes from collisions and dropped packets. Also interesting to note is that EMIT is also able to transmit at very high rates. This is because with very short packet transmission durations, it is less likely for two transmissions to occur simultaneously. This is also evident from the relatively low collisions experienced by

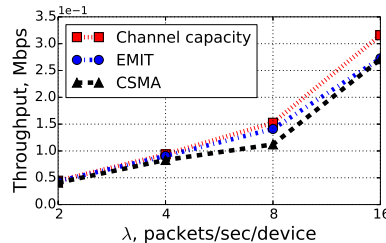


Fig. 8: Throughput of CSMA and highest power setting of EMIT at the USRP testbed.

CSMA. From Figures 7(b) and 8 it can be inferred that EMIT's throughput increases and remains above CSMA's as transmission duration increases and power decreases.

Another point of concern is whether such a paradigm can co-exist with a ubiquitous protocol like Wi-Fi, i.e. **What is the effect of Wi-Fi and EMIT each other?**

G. Co-existence of EMIT and Wi-Fi

1) *Effect of EMIT on Wi-Fi*: In order to observe the effect of EMIT on Wi-Fi, the throughput of a Wi-Fi network is measured with and without EMIT running beside it. For this purpose, commercial Wi-Fi cards at the grid testbed are used to generate traffic using 8 transmitters and 1 receiver, while the USRPs generate EMIT based interference. The Wi-Fi nodes transmit fixed number of packets of the same size at a constant bitrate. It was observed that the packet transmission and reception rate of the Wi-Fi network was unaffected by EMIT for all equal energy settings and λ .

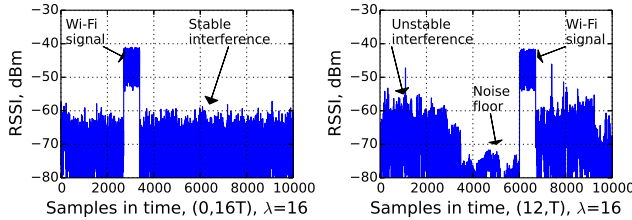


Fig. 9: Effect of EMIT on Wi-Fi. in the USRP testbed

Figure 9 shows the channel observed at a USRP receiver. At low power settings, the interference from multiple EMIT nodes will add up. However the transmission power of the Wi-Fi nodes is much higher than the total interference power generated by EMIT nodes at this setting. When EMIT nodes operate at a high power and short transmission duration settings, they do not occupy the channel often, thereby giving the Wi-Fi nodes ample time to transmit or retransmit packets in case of collisions. This means that Wi-Fi can coexist with EMIT with little to no impact on its throughput.

2) *Effect of Wi-Fi on EMIT*: To examine the effect of Wi-Fi transmissions on interference averaging, a similar setup as earlier is used. Figure 10 shows the amount of increase in variance of interference observed due to Wi-Fi transmission. Variance of interference increases in all cases, but the settings with high λ are affected the less. This is because, for a given amount of Wi-Fi traffic, if the number of EMIT based device transmissions increases, then total energy of the channel increases, and the impact of the Wi-Fi transmissions on the channel variance decreases. The effect of high λ can be achieved by using more devices with a lower λ . Thus denser EMIT based networks can better coexist with Wi-Fi compared to sparser one.

V. RELATED WORK

The current wireless resource allocation strategies used for multi-user interference management can be broadly grouped into two categories: interference avoiding and interference embracing. Interference avoiding strategies,

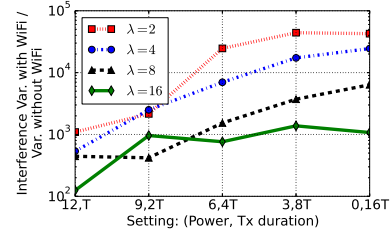


Fig. 10: Increase in variance of interference in the presence and absence of Wi-Fi transmissions.

including slot-based and CSMA-based schemes (e.g. [10], [11], [12]) schedule transmissions depending on the observed interference and load. Such schemes suffer from high and variable delays when used in large network deployments[1]. Interference embracing schemes, like interference alignment and hierarchical cooperation, employ simultaneous transmission with an intricate engineering of multi-user interference. EMIT lies under the interference-embracing paradigm.

Interference alignment was first introduced in [13] and [14]. A different perspective is given in [15] to apply interference alignment in dense wireless networks with ergodic channels between every pair of nodes. The idea of interference alignment is based on exploiting orthogonal parallel resources that vary independently from each other. One can view delay, bandwidth or number of antennas as such resources. To achieve the maximum achievable performance, the amount of resources must scale super-exponentially (see e.g., [13], [16], [17]) with the number of users. As such resources are scarce in IoT, interference alignment is not scalable to large networks.

In [16] and [17], the authors show that delay can be reduced by sacrificing some of the achievable rate, but these methods provide an improvement of a constant factor and do not change the way the delay scales with the number of links in the channel (unless the gain in the rate is completely lost). Furthermore, as the network size grows, achievable rates become *power limited*, rather than *interference limited*. Consequently, in a network spread in a large geographic area, interference alignment cannot be implemented on a global basis⁸.

The idea of clustering and hierarchical cooperation have been proposed in [20], [21], [22] for the MIMO-interference channel and [23] and [24] for the broadcast and the MAC channels, respectively. In those studies, the main purpose of clustering is to achieve optimal capacity scaling in extended networks, whereas the issues of delay and resource usage and coordination overhead have been of secondary importance. Perhaps more importantly, to successfully realize the above schemes, each node requires global channel state information (CSI) in real time. The transmitted symbols over any given channel is a function of

⁸Unlike the dense IoT setting we are targeting, for interference alignment in small-scale networks (with a few transmitter-receiver pairs), there have been some encouraging approaches such as [18] and [19] that address some of the implementation issues.

the state of *all* the channels in the network. In a distributed implementation, this requires a high number of message exchanges between all nodes, in every fading block. The cost of the overhead due to such a message exchange is very high, making the above cooperative approaches non-viable in our envisioned setup of dense IoT.

These paradigms have certain issues that may be problematic for communication in dense IoT networks. First, those that involve careful engineering of interference rely on long-term connectivity and intense local/global coordination among users to achieve a desirable performance. Also, network-level synchronization and information sharing overhead that is already high in existing solutions (see [25], [26]) become even more significant in the presence of such user dynamism, rendering them impractical in their current form. Indeed, in certain cases, such connections may cause existing solutions drive the system to instability as first exposed in [27]. Moreover, the dynamism and delay restrictions of IoT applications break down the traditional time-scale separation assumption between the physical and MAC layers, and user dynamics (in ascending order of time-scales), which underlies the design of most existing communication protocols. Thus, the time scale of user dynamism, coupled with the delay-sensitive nature of intermittently generated data bundles, necessitate the design of low-overhead resource allocation strategies that can operate efficiently under the reversed ordering of the time scale separation.

VI. CONCLUSION

In this paper, we addressed the need for a new MAC paradigm for supporting the communication requirements of emerging IoT networks. These networks will possess non-traditional traffic demands for communicating *small bundles of time-sensitive data* that arrive *densely in space and intermittently in time*. We introduced the principles of EMIT, a multi-user operation strategy that exploits the unique dynamics of IoT networks to shape the global interference through the operation of individual devices. This strategy based on *interference-averaging* stands in contrast to the predominant *interference-avoidance* based MAC solutions, and is well-suited for IoT networks due to its delay and reliability benefits. We performed extensive experimentation to validate and better contrast the benefits of our EMIT paradigm to CSMA-based solutions. We observed that EMIT offers a sharp improvement over CSMA-based solution in supporting IoT networks. This framework is expected to lay the foundations for and inspire new research and development efforts that are better suited for the non-standard characteristics of emerging IoT networks.

Acknowledgement: This work is supported by The National Science Foundation grants CNS-NeTS-1514260, CNS-1054738, CNS-WiFiUS-1456806, CCSS-EARS-1444026, and the Defense Threat Reduction Agency grant HDTRA1-15-1-0003. The work of A. Eryilmaz was supported in part by the Qatar National Research Fund grant NPRP 7-923-2-344.

REFERENCES

- [1] M. Shirvanimoghaddam, Y. Li, M. Dohler, and B. Vucetic, "Probabilistic rateless multiple access for machine-to-machine communication," *arXiv preprint arXiv:1507.02761*, 2015.
- [2] "Anonymized report," <https://sites.google.com/site/infocom16rpf/>.
- [3] N. Sugavanam, C. E. Koksal, and A. Eryilmaz, "Color of interference and joint encoding and medium access in large wireless networks," in *Proceedings of Asilomar*. IEEE Press, 2012.
- [4] F. Kelly, A. Maulloo, and D. Tan, "Rate control for communication networks: shadow prices, proportional fairness and stability," *Journal of the Operational Research Society*, vol. 49, no. 3, March 1998.
- [5] R. Srikant, *The Mathematics of Internet Congestion Control*. Birkhauser, 2004.
- [6] A. Eryilmaz and R. Srikant, "Joint congestion control, routing and mac for stability and fairness in wireless networks," *IEEE Journal on Selected Areas in Communications*, vol. 24, pp. 1514–1524, 2006.
- [7] S. Shakkottai and R. Srikant, *Network Optimization and Control*. Hanover, Massachusetts: Now Publishers, 2008.
- [8] E. Magistretti, O. Gurewitz, and E. W. Knightly, "802.11ec: Collision avoidance without control messages," in *Proceedings of the 18th Annual International Conference on Mobile Computing and Networking*, ser. Mobicom '12. ACM, 2012, pp. 65–76.
- [9] R. U. WINLAB, "Orbit testbed software architecture: Supporting experiments as a service," 2005.
- [10] P. Marbach, A. Eryilmaz, and A. Ozdaglar, "Achievable rate region of csma schedulers in wireless networks with primary interference constraints," in *Proc. IEEE Conference on Decision and Control (CDC'2007)*, New Orleans, Louisiana, December 2007.
- [11] L. Jiang, M. Leconte, J. Ni, R. Srikant, and J. Walrand, "Fast mixing of parallel glauber dynamics and low-delay csma scheduling," in *Proceedings of IEEE INFOCOM Mini-Conference*, 2011.
- [12] A. Sridharan, C. E. Koksal, and E. Uysal-Biyikoglu, "A Greedy Link Scheduler for Wireless Networks with Gaussian Multiple Access and Broadcast Channel," in *Proc. IEEE INFOCOM*, 2010.
- [13] V. R. Cadambe and S. A. Jafar, "Interference alignment and degrees of freedom of the k -user interference channel," *IEEE Transactions on Information Theory*, vol. 54, no. 8, pp. 3425–3441, August 2008.
- [14] M. A. Maddah-Ali, A. S. Motahari, and A. K. Khandani, "Communication over MIMO X Channels: Interference Alignment, Decomposition and Performance Analysis," *IEEE Transactions on Information Theory*, vol. 54, no. 8, pp. 3457–3470, August 2008.
- [15] B. Nazer, M. Gastpar, S. Jafar, and S. Vishwanath, "Ergodic Interference Alignment," in *Proceedings of IEEE International Symposium on Information Theory (ISIT)*, 2009.
- [16] O. Johnson, M. Aldridge, and R. Piechocki, "Delay-rate Tradeoff in Ergodic Interference Alignment," Website, <http://www.arxiv.org/abs/1004.0208>.
- [17] J. C. Koo, W. Wu, and J. T. Gill, "Delay-Rate Tradeoff in Ergodic Interference Alignment in the Gaussian Case," Website, <http://www.arxiv.org/abs/1001.2582>.
- [18] S. Gollakota, S. D. Perli, and D. Katabi, "Interference Alignment and Cancellation," in *Proceedings of ACM SIGCOMM*, 2009.
- [19] L. E. Li, R. Alimi, D. Shen, H. Viswanathan, and Y. R. Yang, "A General Algorithm for Interference Alignment and Cancellation in Wireless Networks," in *Proceedings of IEEE INFOCOM*, 2010.
- [20] A. Ozgur, O. Leveque, and D. N. C. Tse, "Hierarchical Cooperation Achieves Optimal Capacity Scaling in Ad Hoc Wireless Networks," *IEEE Transactions on Information Theory*, vol. 53, 2007.
- [21] U. Niesen, P. Gupta, and D. Shah, "On Capacity Scaling in Arbitrary Wireless Networks," *IEEE Transactions on Information Theory*, vol. 56, pp. 3959–3982, September 2009.
- [22] J. Ghaderi, L. L. Xie, and X. Shen, "Hierarchical Cooperation in Ad Hoc Networks: Optimal Clustering and Achievable Throughput," *IEEE Transactions on Information Theory*, vol. 55, August 2009.
- [23] H. Keshavarz, L. Xie, and R. Mazumdar, "On the optimal number of active receivers in fading broadcast channels," *IEEE Transactions on Information Theory*, vol. 54, pp. 1323–1327, March 2008.
- [24] H. Keshavarz, L. Xie, and R. Mazumdar, "User Capacity Scaling Laws for Fading Multiple-Access Channels," *IEEE Transactions on Wireless Communication*, vol. 8, pp. 4498–4507, September 2009.
- [25] O. El Ayach, S. W. Peters, and R. Heath, "The practical challenges of interference alignment," *Wireless Communications, IEEE*, vol. 20, no. 1, pp. 35–42, 2013.
- [26] P. Verkaik, "Enhancing networking protocols in widely deployed devices," Ph.D. dissertation, Computer Science, University of California, San Diego, 2010.
- [27] P. M. van de Ven, S. C. Borst, and S. Shneer, "Instability of maxweight scheduling algorithms," in *INFOCOM*. IEEE, 2009.