Transmit Only: An Ultra Low Overhead MAC Protocol for Dense Wireless Systems

Bernhard Firner, Yanyong Zhang, Richard Howard, Richard P. Martin,
Narayan Mandayam, Junichiro Fukuyama, Chenren Xu*

Abstract—The number of small wireless devices is rapidly increasing, making the radio channel efficiency in limited geographic areas (individual rooms or buildings) an important metric for MAC protocols. Many of these emerging devices have use-cases that are difficult to satisfy with current hardware solutions and channel access methods; for instance device mobility, small energy reserves, and requirements for low cost and small form factors. However, for most of these applications, such as health care monitoring or sensing, feedback to the radio device is unnecessary and unidirectional communication techniques are not only sufficient, but can also be advantageous. We propose an efficient, reliable technique for unidirectional communication, called Transmit Only (TO), that satisfies these requirements while maintaining packet throughput guarantees and reducing energy consumption. In this paper we will demonstrate the feasibility and performance of this kind of highly asymmetric, transmit-only protocol through theoretical, simulated, and experimental results.

Index Terms—Transmit-Only, Low Power

1 Introduction

The size, cost, and flexibility of current low power radios and microcontrollers is enabling many new computing applications where wireless communication serves as the means to collect data at a finer spatial and temporal resolution than was previously practicable. Although mainstream wireless services, like cellular and WLAN, will continue to grow in adoption and will be well served by capacity enhancements at the physical layer (e.g. MIMO technologies have been responsible for the growth of data rates achieved by LTE and WiMAX), there is another, growing class of wireless-enabled applications, that will not benefit from such physical layer advancements. Densely deployed systems of embedded wireless devices, as proposed in machine-to-machine applications like smart buildings and smart grids [3], [8], [10] or the Internet of Things [6], must operate with a variety of practical constraints that prevent them from benefiting from complex physical layer approaches that enhance capacity and would otherwise support higher deployment densities.

Many communication and networking systems have been built upon a simple assumption: multiple nodes communicating at the same time should be avoided as much as possible because the collisions that arise cannot be dealt with by a receiver. This philosophical assumption is the basis of many medium access control (MAC) protocols.

Regardless of whether one considers time division, frequency division, or carrier sensing, the underlying design objective for these protocols is that transmitters need to avoid collisions before transmission. To avoid collisions, devices either sense the channel or receive a transmission schedule from controllers. Collision avoidance translates into additional hardware requirements for receiver functionality (e.g. channel sensing) causing devices to spend additional energy on functions not directly related to conveying data. These function also cost channel time, resulting in inadequate channel capacity when networks scale to ultra dense deployments.

A dense deployment of embedded wireless devices has no inherent need for feedback to the majority of devices – their sole function is to send data, not to engage in a bidirectional exchange of information. Further, the data being reported by these devices is quite often loss tolerant, so guaranteed delivery of each data packet is not necessary. This is quite different from the design requirements placed on traditional MACs, where data delivery guarantees and bidirectional information exchange were considered vital. Re-examining these requirements leads us to a re-examination of how devices should access the channel: If receiving functions are costly, energy hungry, non-scalable, and unnecessary for guaranteed data delivery or bi-directionality, then if we remove these features can we achieve better performance in dense wireless systems?

- *B.Firne is with the Autonomous Driving Team at Nvidia, Holmdel, NJ . E-mail: {ben.firner}@gmail.com.*
- *Y.Zhang, R.Howard and N. Mandayam are with the Wireless Information Networks Laboratory (WINLAB), Rutgers University, North Brunswick, NJ 08902. E-mail: {yyzhang, reh, narayan}@winlab.rutgers.edu, gv@milvius.com.*
- *R.Moore and R.Martin are with Department of Computer Science, Rutgers University, Piscataway, NJ 08854. E-mail: {romoore, rmartin}@cs.rutgers.edu.*
- *J.Fukuyama is with Applied Research Laboratory, Pennsylvania State University, University Park, State College, PA 16801, United States. E-mail: jxf140@gmail.com.*
- *C.Xu is with the Center for Energy-efficient Computing and Applications, School of EECS, Peking University, Beijing 100871. E-mail: chenren@pku.edu.cn.*
- **Corresponding authors.*

In this paper we will describe and demonstrate the feasibility and performance of a highly asymmetric, transmit-only protocol, called TO. In Section 2 we will discuss past and current research in this area and introduce the idea of the capture effect as a means to achieve increased channel capacity in a small physical area. We will also show that traditional MAC protocols, such as CSMA, do not balance throughput and energy efficiency in dense systems as well as TO, and that TO is allows more efficient use of limited energy and channel resources in dense deployments. In Section 3, we propose a topological approach to maximizing the capture effect with spatial diversity in receiver placement. In Section 4 we will demonstrate the practicality of our approach through experiments in an outdoor space and a complex indoor environment.

2 Overview of Transmit-Only

In this section, we first explain that many MAC protocols are not suitable for dense wireless systems due to their overhead in collision management at the sender's side. We then describe and model our Transmit-Only (TO) protocol and analytically compare it to CSMA.

2.1 Control Overhead in Existing Protocols

The throughput of wireless systems is well-studied. In particular, p-persistent CSMA has received a great deal of attention [13]. In p-persistent CSMA the sender checks whether the channel is idle before transmission. If the channel is idle, the sender transmits a packet with a probability p. By choosing a very small p such as $p = 0.01$, p-persistent CSMA achieves very high throughput even in dense deployments, but at the cost of delayed packet deliveries and poor energy efficiency. A low p means many times an idle channel is detected lead to no packet transmission, costing energy in carrier sensing without any gain. When packets are short compared to the sensing duration, this cost is substantial and has a relatively large impact on lifetime.

Non-persistent CSMA systems will spend less time carrier sensing, but the time spent backing off will result in long packet latencies; if the transmission rate falls below the duty cycle of a device then eventually its transmission queue will fill and packets will be lost. Therefore, though p-persistent and non-persistent CSMA achieve good throughput, they do not balance throughput with energy efficiency and packet latency.

TDMA, CDMA, and FDMA systems all require a number of control packets, and have a variety of weaknesses as transmitter density increases. Control overhead in these systems have a negative impact on power consumption, especially in mobile environments where transmitters frequently encounter new receivers and must negotiate new parameters (time slots, power levels, or frequencies).

2.2 Transmit-Only

These protocols spend a great deal of their energy dealing with or avoiding packet collisions. We point out though, that in some systems collision avoidance can be not only unnecessary, but also very wasteful. For these situations we propose a fundamentally asymmetric architecture call Transmit-Only (TO), where the transmitters are only able to transmit, eliminating overhead associated with all interference avoidance techniques or channel feedback. In a TO network, we deploy multiple receivers and ensure that each transmitter is within a single hop of one or more receivers. The receivers work on the same channel and exploit the *capture effect* to reduce the effective contention between transmitters.

The capture effect has received some attention in the literature. For instance, Takagi and Kleinrock [12] considered the benefits of the capture effect when multiple receivers are present in multi-hop networks. They found that the ALOHA protocol benefits more from the capture effect than 1-persistent CSMA. More importantly, they found that with "perfect capture" where during any collision the one packet with the strongest signal power will be received regardless of its relative power compared to the other packets, ALOHA could outperform CSMA, while ALOHA with 1.5dB capture (i.e., packet capture occurs when the relative power is 1.5dB or greater) performed comparably to CSMA. These comparison results of CSMA and ALOHA hint that the capture effect can effectively reduce data loss even in the presence of collisions. In this paper, we thoroughly study methods that can maximize the capture gains.

2.3 Comparison of TO and CSMA

We first provide an analytical comparison of p-persistent CSMA and TO (TO does not have any feedback from the receiver as in ALOHA). This comparison allows us to show the possible benefits of having multiple receivers under a theoretically ideal situation and will motivate the remainder of this paper. To begin our analysis we will assume that we have "perfect" capture - the strongest signal can always be captured. We will discuss how to maximize potential capture gains in real-world situations in Section 3. In Section 4, we show that our experimental results with realistic radio parameters are in agreement with our analytical results.

In TO the probability of two transmitters having a packet collision depends upon the duty cycle of transmitters. We will consider a simple case where traffic is periodic with interval τ and durations δ (a duty cycle of δ/τ). TO is unslotted, so a collision occurs when any overlap occurs, or $2\delta/\tau$. The probability of two transmitters colliding is

$$P_{2-way-collision} = \frac{2\delta}{\tau}. \quad (1)$$

With N transmitters, a transmitter's packet is received if no collisions occur, the probability of which is

$$P_{succ} = (1.0 - \frac{2\delta}{\tau})^{N-1}. \quad (2)$$

In the presence of contention, a transmitter may collide with any of the $N-1$ other transmitters, but due to the capture effect, its packet can be correctly received if it collides with transmitters that have lower signal strengths at a receiver. Let us assume that, even with the capture effect at the receivers, a transmitter's packet cannot be received when it collides with any of the C transmitters (a subset of $N-1$ other transmitters). We say that this transmitter's *contention level* is C, and the packet success probability is thus:

$$P_{succ} = 1.0 - (1.0 - \frac{2\delta}{\tau})^C. \quad (3)$$

This equation gives a good intuition for packet success, but it does not include details of capture when there are multiple receivers or when it is difficult to determine the contention level of each transmitter. In this case we must model the system with a more exact equation that takes into account the probability of a transmitter's packet being captured in a collision involving any number of other transmitters, from 1 to $N-1$. The probability of packet loss from a collision is simply a binomial random variable with the addition of the capture probability with each collision magnitude.

$$P_{loss} = \sum_{i=1}^{N-1} \left(\frac{2\delta}{\tau}\right)^i \left(1 - \frac{2\delta}{\tau}\right)^{N-i-1} \binom{N-1}{i} (1 - P_{capture}). \quad (4)$$

In the case of *perfect capture*, we assume that during any packet collision with any number of packets a receiver will always correctly decode one of the packets. In this case the probability of any transmitter involved in an n-way collision having its packet captured is simply $1/n$. Given n transmitters and r receivers the probability of a particular transmitter not having its packet captured is

$$1 - P_{perfect-capture}(n, r) = (1 - 1/n)^r. \quad (5)$$

In non-perfect capture case we need to compute both the probability that during a collision a particular transmitter is the strongest transmitter and also that the strongest transmitter is actually captured over all other transmitters. First, we will let K be the likelihood that one packet is captured over another and assume that probability of capture at different receivers is independent. The probability that the one of the packets in a collision is captured over all of the others is the probabilty of $n-1$ captures taking place, or K^{n-1}. Combining this with the probability of a particular transmitter being that strongest signal and considering the possibility of capture at any of r receivers when n transmitters collide we find

$$1 - P_{capture}(n, r) = \left(1 - \frac{K^{n-1}}{n}\right)^r. \quad (6)$$

In related work the capture probability has been successfully modeled as a relationship between the capture threshold of the hardware being used, Δ, and the propagation coefficient of an environment, α, assuming a propagation model of $1/r^\alpha$ [2]. Given a uniformly random deployment of transmitters radiating from the location of a receiver, the probability of one particular transmitter having its packet captured is

$$K = \frac{10^{-\Delta/10\alpha}}{2}. \quad (7)$$

Using Equations 4 and 5 we can calculate the theoretical gains of the perfect capture effect with multiple receivers, as shown in Fig. 1. In Fig. 1, we use *offered packets per packet interval* as the offered load. From this result we can see that the capture effect could greatly increase throughput in TO compared to a single receiver ALOHA system without capture.

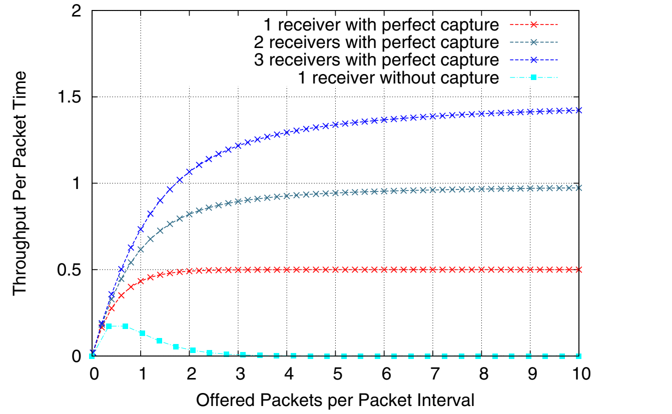

Fig. 1. ALOHA without capture has a low packet throughput, but when we consider the capture gains from multiple receivers very high packet throughputs are achievable.

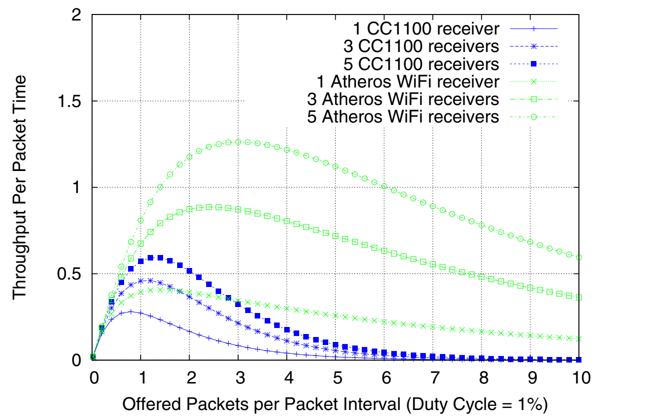

Fig. 2. The hardware's capture threshold has a large effect upon performance as shown in a comparison of the expected throughput of TO with different numbers of CC1100 radios Atheros WiFi cards. $\alpha = 3$ is assumed to find the capture rate from Equation 7.

Using Equations 6 and 7 we can also estimate performance under non-perfect capture conditions, for instance by using the 6dB [2] capture threshold of a CC1100 radio or the 1dB [7] threshold of some Atheros WiFi cards. A very important difference between perfect and non-perfect capture is that performance degrades as the offered load increases in the non-perfect case. The rate of this change in performance is actually a factor of the packet duration and the interval between packet (the duty cycle), as can be seen by comparing Figs. 2 and 3.

To compare TO with CSMA, we adopt the radio efficiency analysis performed by Ramachandran and Roy [9]. We first show the throughput values for different p-persistent CSMA and TO protocols in Fig. 4. The results clearly show that when we have 3 or more receivers the throughput of TO with perfect capture is higher than the CSMA protocols. In non-perfect cases the same result can be achieved, albeit with more receivers.

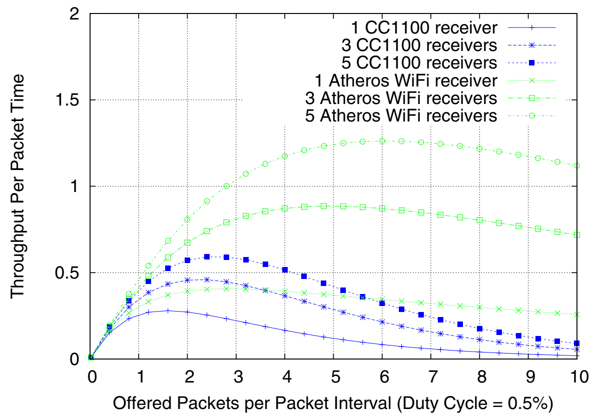

Fig. 3. In the non-perfect capture case the ratio of packet duration to packet interval, the duty cycle, affects when the system reaches peak performance. Longer packets cause the system to have smaller "good" ranges.

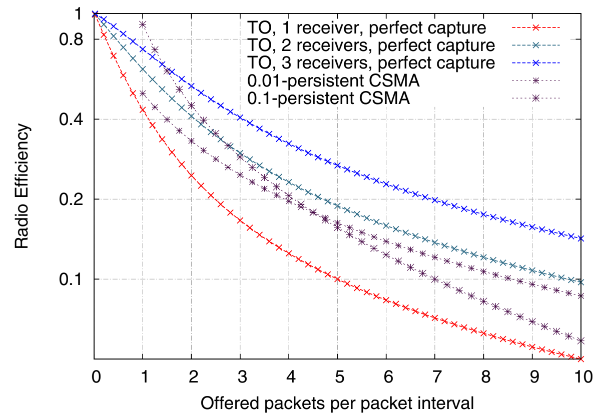

Fig. 5. Although 0.01-persistent CSMA may give good throughput, it is not energy efficient. TO is able to achieve both good throughput and good energy consumption.

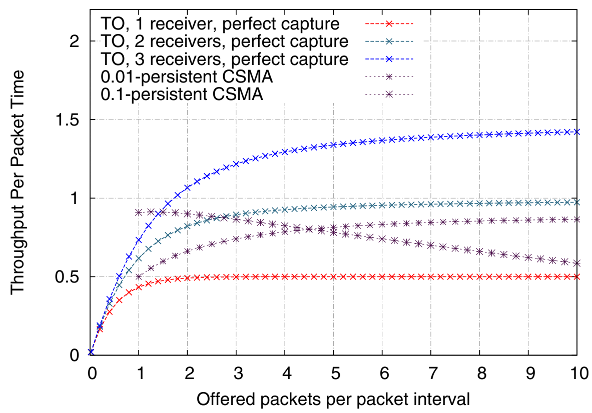

Fig. 4. A comparison of throughput for TO and p-persistent CSMA with different p-values.

We show the corresponding radio efficiency results in Fig. 5. Radio efficiency is the ratio between the time spent in successfully transmitting data packets and the total radio time. The energy requirements of the radio usually dominate the energy consumption of wireless devices, so this may be the single most important factor for device lifetime. Again, in TO systems having 3 or more receivers can greatly improve the radio efficiency, and thus device lifetimes. Specifically, when the offered packets per packet is around 3, TO with 4 receivers triples the radio efficiency compared to both CSMA protocols.

3 Maximizing the Capture Gains

The capture effect in wireless radios increases throughput by reducing contention between transmitters. The locations of receivers in a wireless network, however, determine when the capture effect occurs, and thus the total amount of contention. By predicting the occurrence of the capture effect given the receiver locations, we can also choose multiple receiver locations that minimize transmitter contention and thus maximize throughput in a TO system. This technique can also be used to predict the number of receivers required to achieve a desired level of packet throughput in a deployment.

3.1 Quantifying the Capture Gain

Section 2 introduced Equations 4, 5, 6, and 7 to describe collision losses when there are multiple transmitters and the capture effect is considered. If we model the contention between transmitters directly we can simply apply Equation 3 to estimate a transmitter's packet loss, as long as the probability of collisions involving many transmitters is low. This gives us a single metric to concentrate on for optimization.

Previous work has successfully described capture success as a function of a capture threshold, Δ (which is hardware dependent), and the propagation coefficient of an environment, α, assuming a propagation model of $1/r^{\alpha}$ [2]. We will call β the minimum ratio of the distance from a far transmitter, d_1, to the distance of a near transmitter whose packet we want to capture, d_2. We describe β as

$$\begin{aligned} (d_2/d_1) &\geq 10^{\Delta dB/10\alpha} \\ \beta &\geq 10^{\Delta dB/10\alpha}. \end{aligned} \quad (8)$$

With knowledge of β we can choose the best receiver placement to achieve high packet capture rates and achieve good fairness in the system.

3.2 Optimal Receiver Placement

We now formally define the optimal receiver placement problem. Consider two transmitters located at $t_1, t_2 \in \mathbb{R}^2$, and a receiver located at $r \in \mathbb{R}^2$. For the sake of simplicity, we will use t_1, t_2 and r as both their locations and identities. In case of a packet collision between t_1 and t_2, the signal from t_1 can be captured by r if and only if

$$\|r - t_1\| \leq \beta \|r - t_2\|, \quad (9)$$

where $\|\cdot\|$ is the Euclidian norm of a vector in $\mathbb{R}^2$, and $\beta \in (0, 1)$ describes the relative distance difference for the capture effect to take place. In this case, we say that the ordered transmitter pair (t_1, t_2) is successfully *captured*

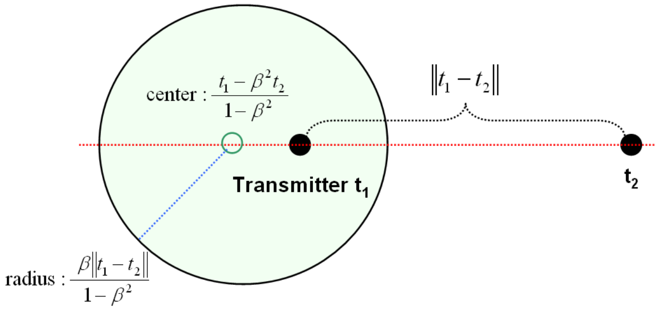

Fig. 6. The capture disk of transmitter pair (t_1, t_2). A receiver placed inside of the disk will capture t_1's packets when they collide with t_2's packets.

by r. Our goal is to *find m receiver locations that maximize the captured transmitter pairs and thus minimizes the average contention*. Formally, our *receiver embedding problem* is defined as follows:

Given: locations $t_1, t_2, \ldots, t_n \in \mathbb{R}^2$ of n transmitters and the number of receivers, m.

Find: m receiver locations $r_1, r_2, \ldots, r_m \in \mathbb{R}^2$ such that the number of captured transmitter pairs is maximum.

3.2.1 F Approximation for Receiver Placement

To solve the optimal problem, we will first redefine the continuous optimization problem as a discrete problem with a finite number of points in $\mathbb{R}^2$. Let us start this process by finding a necessary and sufficient condition to (9).

$$\begin{aligned} & \|r - t_1\|^2 \leq \beta^2 \|r - t_2\|^2 \\ \Leftrightarrow \quad & (1-\beta^2)\|r\|^2 - 2r \cdot (t_1 - \beta^2 t_2) \leq -\|t_1\|^2 + \beta^2 \|t_2\|^2 , \end{aligned}$$

where $\cdot$ takes the inner product of two vectors. It is then further transformed into

$$\left\| r - \frac{t_1 - \beta^2 t_2}{1-\beta^2} \right\| \leq \frac{\beta}{1-\beta^2} \|t_1 - t_2\| . \quad (10)$$

The region of possible receiver location r such that (t_1, t_2) is captured by r is a disk centered at $\frac{t_1-\beta^2 t_2}{1-\beta^2}$ and with radius $\frac{\beta}{1-\beta^2}\|t_1 - t_2\|$. This *capture disk of* (t_1, t_2) is illustrated in Fig. 6. This means that (t_1, t_2) is captured if and only if its capture disk contains a receiver.

Now we can formulate a discrete version of the receiver embedding problem with two steps. First, given n transmitters, we compute the capture disks of all the ordered transmitter pairs $(t_i, t_{i'})$ $(i \neq i')$. Second, since all points inside of a disk have the same utility for capturing collisions we only need to count each disk once, as well as any points that lie inside of multiple disks. The possible *solution points* are thus the center of each capture disk and the intersection point between the boundary circles of any two intersecting capture disks.

The following lemma shows the above discrete problem is equivalent to the original receiver embedding problem.

Lemma 3.1. *For any given transmitters $t_1, t_2, \ldots, t_n$ for the receiver embedding problem, there exist optimal m receiver locations $r_1, r_2, \ldots, r_m$ such that every r_j ($1 \leq j \leq m$) is a solution point.*

Algorithm F-EMBED
Inputs:
1. n transmitter locations $t_1, t_2, \ldots, t_n \in \mathbb{R}^2$.
2. Positive integer m.

Output: m receiver locations $r_1, r_2, \ldots, r_m \in \mathbb{R}^2$.
begin
1. Compute the center and radius of the capture disk of every ordered transmission pair $(t_1, t_2), (t_1, t_3), \ldots, (t_n, t_{n-2}), (t_n, t_{n-1})$;
2. Compute all the solution points, *i.e.*, the centers and intersections of capture disks;
3. Construct a bipartite graph $G = (S, T, E)$ as follows;
   3-1. $S = \{s_1, s_2, \ldots, s_l\}$ is the set of solution points;
   3-2. $T = ((t_1, t_2), \ldots, (t_n, t_{n-1}))$ is the set of ordered transmitter pairs;
   3-3. E is the edge set such that $(s_j, (t_i, t_{i'})) \in E \Leftrightarrow$ the capture disk of $(t_i, t_{i'})$ contains the solution point s_j;
4. **for** $k = 1$ to m **do**
   4-1. Find a solution point $s_j \in S$ connected to a maximum number of transmitter pairs $(t_i, t_{i'})$;
   4-2. Let the k^{th} receiver location r_k be this s_j;
   4-3. Delete s_j and $(t_i, t_{i'})$ connected to s_j from G;
5. **end for**

end

Fig. 7. Our F-EMBED Algorithm to Find m Receiver Locations

Proof. Suppose that an optimal receiver location r_j captures p transmitter pairs. It means that r_j belongs to the intersection I of p capture disks. Any such region I contains an intersection between the boundary circles of two capture disks, or the center of a capture disk. Let $r'_j \in \mathbb{R}^2$ be such a solution point.

Given any optimal set $\{r_1, r_2, \ldots, r_j\}$ of receiver locations, construct another by replacing each r_j by the above solution point r'_j. We have a solution consisting of only solution points. This proves the lemma. □

In the discrete version we need only examine a finite set of solution points. Assuming finite digits are used to represent a vector in $\mathbb{R}^2$, this discrete version is similar to classical **NP**-complete problems such as vertex cover and subset sum. We thus believe that both the original and discrete versions of the receiver embedding problem are **NP**-hard and will only seek an approximation algorithm of constant factor to the optimal solution.

We will present a 2-approximation algorithm for the receiver embedding problem, *i.e.*, it always returns a solution that captures at least half as many collisions as the optimum solution. Our algorithm, called F-EMBED, has four steps:

- In step 1, we compute the capture disks for all the ordered transmitter pairs.
- In step 2, we compute the solution points.
- In step 3, we construct a bipartite graph G such that solution point s_j is connected to transmitter pair $(t_i, t_{i'})$ if and only if s_j is contained in the capture disk of $(t_i, t_{i'})$.
- In step 4, we go through m iterations and in each iteration we pick the solution point that has the maximum number of edges in the graph and remove the solution point as well as its edges.

A walk through of the algorithm using a 3-transmitter, 2-receiver example is illustrated in Fig. 8. The quality of

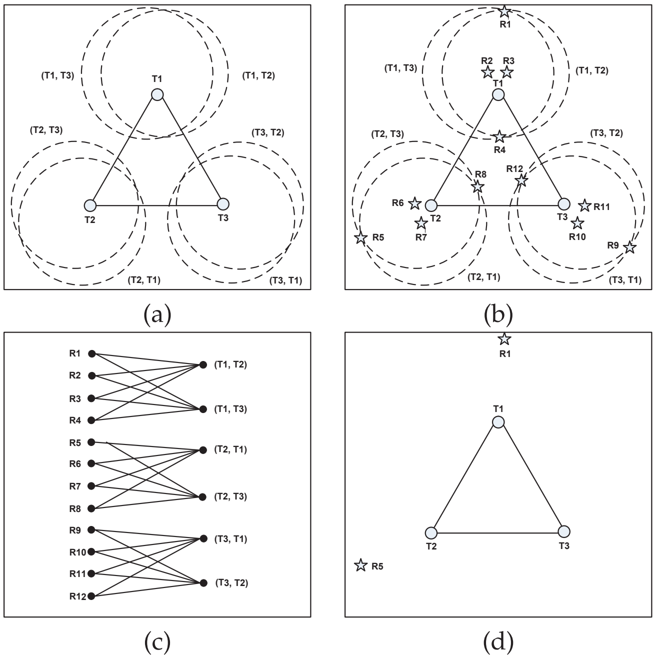

Fig. 8. Walking through the proposed receiver embedding solution with a 3-transmitter, 2-receiver example. We calculate the capture disks for all six ordered transmitter pairs in (a), and generate their solution points in (b). We show the bipartite graph between solution points and transmission pairs in (c). (d) shows the resulting receiver locations. Notice that there are multiple solutions with equivalent performance.

the solution that this algorithm achieves is an important result though, so we confirm in the following theorem that its approximation factor is 2.

Theorem 3.2. F-EMBED *is a 2-approximation algorithm for the receiver embedding problem.*

Proof. Given transmitter locations $t_1, t_2, \ldots, t_n$ and an integer m for the embedding problem, let o_S and c_S denote the optimal and F-EMBED receiver locations, and o_T and c_T denote the ordered transmitter pairs captured by o_S and c_S respectively.

By Lemma 3.1, we assume that the optimal solution o_S consists of solution points only. This means that

$$o_S, c_S \subseteq S = \{s_1, s_2, \ldots, s_l\}$$
$$\text{and} \quad o_T, c_T \subseteq T = \{(t_1, t_2), (t_1, t_3), \ldots, (t_n, t_{n-1})\}.$$

Here S and T are the vertex sets of the bipartite graph $G = (S, T, E)$ constructed by Step 3 of F-EMBED.

We prove by induction on m that

$$|c_T| \geq \frac{1}{2} |o_T|. \tag{11}$$

The base case $m = 1$ is clearly true. Let us assume true for $m - 1$ and prove true for m.

Without loss of generality, assume that the solution point chosen by Step 1 is $s_l \in S$. Let $T_l \subseteq T$ be the (set of) ordered transmitter pairs captured by s_l. Step 3 removes them from G, so the remaining solution points are

$$S' = \{s_1, s_2, \ldots, s_{l-1}\} \subset S,$$

and transmitter pairs

$$T' = T - T_l.$$

If $s_l \in o_S$, then (11) holds true. By induction hypothesis, F-EMBED returns c_S such that $|c_T \cap T'| \geq \frac{1}{2} |o_T \cap T'|$. Thus

$$|c_T| = |c_T \cap T'| + |T_l| > \frac{1}{2} (|o_T \cap T'| + |T_l|) = \frac{1}{2} |o_T|,$$

proving (11).

So we assume $s_l \notin o_S$. The optimal solution o_S chooses m solution points in $S' = S - \{s_l\}$. Let s_{l-1} denote an element in $o_S \subset S'$, and T_{l-1} denote the transmitter pairs captured by s_{l-1}. We have

$$\begin{aligned} |c_T| &= |c_T \cap T'| + |T_l| \geq \frac{1}{2} (|o_T \cap T'| - |T_{l-1}|) + |T_l| \\ &\geq \frac{1}{2} (|o_T| - |T_l| - |T_{l-1}|) + |T_l| \\ &\geq \frac{1}{2} (|o_T| - 2 |T_l|) + |T_l| \geq \frac{1}{2} |o_T|. \end{aligned}$$

This completes the proof of the induction step. □

4 DEPLOYMENT AND EVALUATION

We performed two real world experiments to validate our models of TO, to illustrate the utility of the ADAPTIVE-EMBED placement algorithm, and to demonstrate the performance of TO under high node density. We will use the first experiment to verify that our simplified model of TO agrees with actual measurements. In the second experiment we will use a structured transmitter topology to show that the ADAPTIVE-EMBED placement can achieve a better throughput than the naive placement. Using the same experiment, we also show that TO can achieve very good throughput; with an offered load of 100% TO with three receivers achieves 92.1% throughput on a single channel.

The experiments were performed in a 10 by 10 meter area of an outdoor field located over 50m from the nearest building (as shown in Fig. 9(a)). In the experiments we used 500 custom-made programmable sensor node[1]. The over the air packets contained 4 bytes of preamble and 4 bytes of synchronization bits, meaning that a minimum packet of a single byte would still include those 8 overhead bytes. Different packet lengths were used in the two experiments and the overhead bytes are counted in the over the air packet durations.

4.1 Model Validation

The purpose of the first experiment is to show that our theoretical models match real world deployments at scale. Each of the 500 deployed transmitters broadcast a 384 μsecond packet once a second. To vary the offered load, transmitters were added into the system 100 at a time, starting from 200 transmitters and ending at 500.

Transmitters were deployed in a uniform random 10m by 10m square as shown in Fig. 9(b). Transmitters were bundled into groups to make the experiment tractable at scale. Fig. 9(a) shows a map of the experiment for the 500 transmitter and 3 receiver case with the ADAPTIVE-EMBED placement.

1. Node details are not included here to follow anonymization guidelines, but will be included in a final version.

Fig. 9. High density experiments in an open area outdoors. Fig. (a) shows the transmitter/receiver locations in a 10 by 10 meter area. Bundles of transmitters are shown as black triangles (shown in detail in (b)) and receivers are shown as white triangles (shown in detail in (c)).

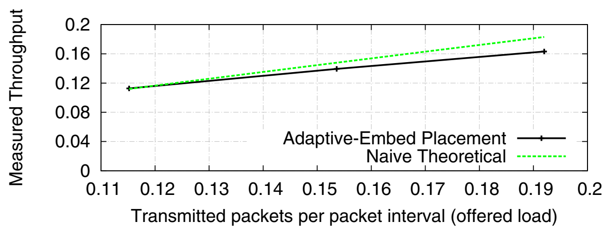

Fig. 10. The throughput of the outdoor high-density experiment is close to our theoretical calculation. In topologies with no underlying structure, such as a uniform grid, receiver in a simple, symmetrical pattern are nearly as effective as those placed with ADAPTIVE-EMBED.

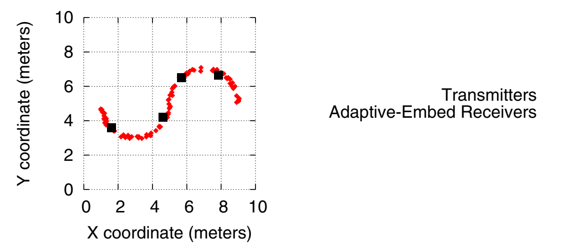

Fig. 11. The topology of the structured, high-load deployment. The Adaptive-EMBED algorithm placed four of the receivers, while four others were evenly spaced for good signal strength.

Fig. 10 shows the the achieved throughput as we scale the number of transmitters. The predictions of our theoretical model are also compared to the measured results of a naive receiver placement (three receivers placed at vertices of an equilateral triangle within the square) and the ADAPTIVE-EMBED placement. The model estimates packet loss by using the average transmitter contention and Equation 3. The results show that even at 500 transmitters, our throughput prediction agrees reasonably with the actual measured value. Also of importantance is that for uniformly placed topologies, simple uniform placement of receivers, i.e., an equilateral triangle, is close to the ADAPTIVE-EMBED placement so a complex placement approach may in unnecessary in many real-world situations.

Fig. 10 also shows the effectiveness of the multiple receiver strategy. As the number of transmitters increases from 200 to 500 the offered load goes from 7.68% to 19.2%. At 200 transmitters the naive placement still received 99.4% of packets, achieving a throughput of 7.63%. At 500 transmitters the ADAPTIVE-EMBED placement lost 15% of packets, achieving a throughput of 16.3%, while the naive placement lost 17% of packets, achieving a throughput of 15.9%.

We note that we do not expect the ADAPTIVE-EMBED placement to significantly outperform the naive placement in this trial. The random uniform random pattern has no underlying structure for ADAPTIVE-EMBED to exploit. The offered load is also rather modest, so the collision rate is not very high. To show how the ADAPTIVE-EMBED placement can improve upon a simple receiver placement strategy the next experiment will use a structured topology for transmitters rather than a uniform random scattering.

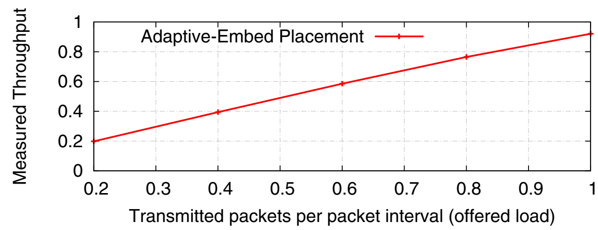

Fig. 12. The throughput of the outdoor experiment with transmitters deployed in a sine wave pattern. Deploying three receivers according to the ADAPTIVE-EMBED algorithm decreased packet loss by more than 35% compared to deploying them to a typical strategy of maximizing RSS values.

We also increase the offered load up to 100% to show how the multiple receiver strategy used in TO can support even higher levels of traffic.

4.2 High Load in a Structured Topology

For our second experiment we increased packet durations to 1ms and increased transmission rates to twice per second. The offered load ranged from 20% with 200 active transmitters to 100% with 500 active transmitters. This time we bundled the transmitters in groups of five and deployed them in a uniform random distribution along a sine-wave shaped line creating a situation that we believe is similar to radios deployed along existing geographical of manufactured paths of interest, such as rivers, roads, and borders. An illustration of the deployment is in Fig. 11.

At low levels of offered load the performance of the simple receiver placement and the ADAPTIVE-EMBED placement were very similar, as shown in Fig. 12. However, as the offered load increased and packet collisions became more probable the ADAPTIVE-EMBED placement begins to perform better. While the simple placement only achieves 87.9% throughput with an offered load of 100%, the ADAPTIVE-EMBED placement reduces packet loss by 35% compared to the simple placement, and achieves 92.1% throughput.

5 Related Work

Low power wireless protocols have received a great deal of attention in the research community, and the works we discuss below are a small set of examples to provide some context. The major goal of most of these protocols is to reduce the amount of time a transmitters radio is on, called its radio duty cycle.

WiseMAC is an example of a MAC protocol that achieves very low duty cycles [1]. In WiseMAC, transmission from a device to an access point (AP) consists of just sensing the channel, transmitting a packet, and receiving an ACK packet. WiseMAC reduces the node duty cycle to 1-2%, which is a significant improvement over previous protocols. SCP-MAC [15] extends this with an adaptive polling frequency that adjusts how often it polls the channel to achieve more efficient performance with low collection duty cycles. However, as data transmission duty cycle increase SCP-MAC must also increase channel polling to the same rate as other protocols.

Some MAC protocols were written to perform well when transmitting to a single sink, similar to the way that TO operates. MAC protocols like Crankshaft [4] were designed to reduce energy consumption through further synchronization to avoid contention. Y-MAC [5] goes one step further and divides traffic across multiple channels through a channel hopping scheme with a single control channel for coordination and TDMA for collision avoidance with contention windows and low-power listening in time slots. Asynchronous protocols like RI-MAC [11] remove the need for synchronization except when there is data to send, while still accommodating high traffic loads. However, the energy consumed searching for a rendezvous time for an intended receiver could be very high. Another similar protocol called EM-MAC [14] uses a predictive mechanism to reduce energy overhead by predicting time slots in advance. Individual transmissions between a static link are made very efficient, but this comes at the cost of a long and energy expensive setup phase, making EM-MAC unsuitable for dynamic deployments with mobile devices.

6 Conclusions and Future Work

In this work we demonstrated the advantages of TO for use in real wireless networks. We demonstrated that with careful placement of receivers, we can leverage the capture effect to recover colliding packets and developed analytic models to approximate optimal receiver placement and predict their performance. Finally, we validated our model via experiments scaling up to 500 transmitters.

Both our models and experiments are in close agreement. These support the feasibility of TO in real deployments because our models and experiments show that the number of receivers scales linearly with the number of transmitters, but that rate of the increase is very modest. With showed that with only five receivers, contention in a network can be reduced by 90%, although we experience diminishing returns as the number of receivers are increased.

References

[1] A. El-Hoiydi and J.-D. Decotignie. Wisemac: an ultra low power mac protocol for the downlink of infrastructure wireless sensor networks. In *Computers and Communications, 2004. Proceedings. ISCC 2004. Ninth International Symposium on*, volume 1, pages 244 – 251 Vol.1, june-1 july 2004.

[2] B. Firner, C. Xu, R. Howard, and Y. Zhang. Multiple receiver strategies for minimizing packet loss in dense sensor networks. In *Proceedings of the eleventh ACM international symposium on Mobile ad hoc networking and computing*, MobiHoc '10, pages 211–220, New York, NY, USA, 2010. ACM.

[3] V. Gungor, B. Lu, and G. Hancke. Opportunities and challenges of wireless sensor networks in smart grid. *Industrial Electronics, IEEE Transactions on*, 57(10):3557 –3564, oct. 2010.

[4] G. P. Halkes and K. G. Langendoen. Crankshaft: an energy-efficient mac-protocol for dense wireless sensor networks. In *Proceedings of the 4th European conference on Wireless sensor networks*, EWSN'07, pages 228–244, Berlin, Heidelberg, 2007. Springer-Verlag.

[5] Y. Kim, H. Shin, and H. Cha. Y-mac: An energy-efficient multi-channel mac protocol for dense wireless sensor networks. In *IPSN '08: Proceedings of the 7th international conference on Information processing in sensor networks*, pages 53–63, Washington, DC, USA, 2008. IEEE Computer Society.

[6] M. Kranz, P. Holleis, and A. Schmidt. Embedded interaction: Interacting with the internet of things. *Internet Computing, IEEE*, 14(2):46 –53, march-april 2010.

[7] J. Lee, W. Kim, S.-J. Lee, D. Jo, J. Ryu, T. Kwon, and Y. Choi. An experimental study on the capture effect in 802.11a networks. In *WinTECH '07: Proceedings of the second ACM international workshop on Wireless network testbeds, experimental evaluation and characterization*, pages 19–26, New York, NY, USA, 2007. ACM.

[8] J. Lu, T. Sookoor, V. Srinivasan, G. Gao, B. Holben, J. Stankovic, E. Field, and K. Whitehouse. The smart thermostat: using occupancy sensors to save energy in homes. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, pages 211–224, 2010.

[9] I. Ramachandran and S. Roy. Analysis of throughput and energy efficiency of p-persistent csma with imperfect carrier sensing. In *Proceedings of the Global Telecommunications Conference (GLOBECOM)*, page 5, St. Louis, MO, USA, Nov. 2005.

[10] A. Roy, S. Das Bhaumik, A. Bhattacharya, K. Basu, D. Cook, and S. Das. Location aware resource management in smart homes. In *Proceedings of the First IEEE International Conference on Pervasive Computing and Communications*, pages 481 – 488, march 2003.

[11] Y. Sun, O. Gurewitz, and D. B. Johnson. Ri-mac: A receiver-initiated asynchronous duty cycle mac protocol for dynamic traffic loads in wireless sensor networks. In *Proceedings of the 6th ACM conference on Embedded network sensor systems*, SenSys '08, pages 1–14, New York, NY, USA, 2008. ACM.

[12] H. Takagi and L. Kleinrock. Optimal transmission ranges for randomly distributed packet radio terminals. *Communications, IEEE Transactions on*, 32(3):246 – 257, mar 1984.

[13] H. Takagi and L. Kleinrock. Throughput analysis for persistent csma systems. *Communications, IEEE Transactions on*, 33(7):627 – 638, jul 1985.

[14] L. Tang, Y. Sun, O. Gurewitz, and D. B. Johnson. EM-MAC: A Dynamic Multichannel Energy-Efficient MAC Protocol for Wireless Sensor Networks. In *The Twelfth ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc11)*, May 2011.

[15] W. Ye, F. Silva, and J. Heidemann. Ultra-low duty cycle mac with scheduled channel polling. In *SenSys '06: Proceedings of the 4th international conference on Embedded networked sensor systems*, pages 321–334, New York, NY, USA, 2006. ACM.