ART: Adaptive fRequency-Temporal co-existing of ZigBee and WiFi

Feng Li, Jun Luo, Gaotao Shi, Ying He

Abstract—Recent large-scale deployments of wireless sensor networks have posed a high demand on network throughput, forcing all (discrete) orthogonal ZigBee channels to be exploited to enhance transmission parallelism. However, the interference from widely deployed WiFi networks has severely jeopardized the usability of these discrete ZigBee channels, while the existing CSMA-based ZigBee MAC is too conservative to utilize each channel temporally. In this paper, we propose ART (*Adaptive fRequency-Temporal co-existing*) as a framework consisting of two components: FAVOR (*Frequency Allocation for Versatile Occupancy of spectRum*) and P-CSMA (*Probabilistic CSMA*), to improve the co-existence between ZigBee and WiFi in both frequency and temporal perspectives. On one hand, FAVOR allocates continuous (center) frequencies to nodes/links in a near-optimal manner, by innovatively converting the problem into a spatial tessellation problem in a unified frequency-spatial space. This allows ART to fully exploit the "frequency white space" left out by WiFi. On the other hand, ART employs P-CSMA to opportunistically tune the use of CSMA for leveraging the "temporal white space" of WiFi interference, according to real-time assessment of transmission quality. We implement ART in MicaZ platforms, and our extensive experiments strongly demonstrate the efficacy of ART in enhancing both throughput and transmission quality.

Index Terms—Continuous frequency allocation, ZigBee, WiFi, CSMA, spatial tessellation.

1 Introduction

Wireless Sensor Networks (WSNs) have been attracting great interests in both academic and industrial communities [5]. Since recent applications, involving both static [6] and semi-mobile [11] indoor scenarios, have posed increasing demands in network performance mainly due to their very large density, the multi-channel capability of ZigBee radios (e.g., CC2420 of MicaZ mote [1]) should be exploited to enable concurrent transmissions [19], [21].

Unfortunately, multi-channel WSNs face the contentions from other networks that also operate in the 2.4 GHz ISM band, e.g., the pervasive deployments of WiFi hotspots [9], [10], [14]. On the one hand, as WiFi radios have a much higher power than ZigBee ones while their frequency bands heavily overlap (see Fig. 1), the WiFi interference may substantially confine the availability of the orthogonal channels specified in ZigBee standard [2] and hence degrade transmission throughput and quality (measured by *Packet Reception Ratio*, or PRR). On the other hand, the most popular mechanism to resolve interference, *Carrier Sense Multiple Access* (CSMA), sacrifices throughput for improving PRR: it simply uses conservative temporal spacing to avoid transmission collisions. These observations have motivated us to re-visit the co-existence issue of ZigBee and WiFi, aiming to ensure both transmission throughput and quality for multi-channel WSNs.

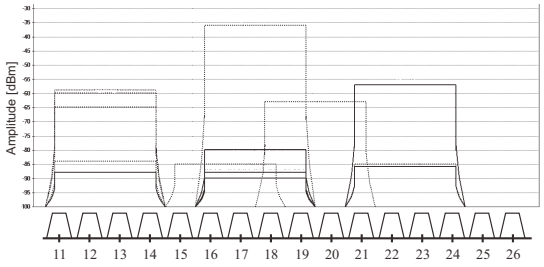

Fig. 1. We overlap the WiFi signal strength observed through inSSIDer in our research center with the 16 channels of ZigBee. Apart from channels 25 and 26, all others are subject to multiple WiFi interferences.

To improve the performance of ZigBee networks under WiFi interference, one needs to maximize the utilization of limited spectrum resource. Inspired by the fact that interference attenuates with distance, our experiments demonstrate it is feasible to allocate continuous center frequencies (thus overlapped non-orthogonal channels) to nodes/links, with respect to their distances to avoid resulting in ZigBee interference. This would be much more efficient than conventional graph coloring approach [16] where discrete channels are allocated based on the concept of discrete graph distance.

While the continuous spectrum utilization method enables us to improve transmission parallelism with limited ZigBee spectrum, a collision avoidance mechanism is still necessary in case of inevitable WiFi interference. Our experiments demonstrate, blindly enabling CSMA regardless of the intensity of WiFi interference offers a high PRR but drastically reducing the throughput. Given the temporal variations exhibited by WiFi interference [10], an adaptive

- *This work is partially supported by The Fundamental Research Funds of Shandong University (No. 11150076614011). Preliminary results were presented in Proceedings of the 14th ACM MobiHoc, 2013 [13].*
- *Feng Li is with School of Computer Science and Technology, Shandong University, China. Email: fli@sdu.edu.cn.*
- *Jun Luo, and Ying He are with School of Computer Science and Engineering, Nanyang Technological University, Singapore. E-mail: {junluo, yhe}@ntu.edu.sg.*
- *Gaotao Shi is with School of Computer Science and Technology, Tianjin University, China. E-mail: shgt@tju.edu.cn.*

MAC protocol that opportunistically tunes the use of CSMA according to interference intensity could achieve both high throughput and high PRR.

Driven by the aforementioned two observations, we propose ART (*Adaptive fRequency-Temporal co-existing*) as a framework to improve the co-existence status between ZigBee and WiFi in a joint frequency-temporal manner. Essentially, ART has FAVOR (*Frequency Allocation for Versatile Occupancy of spectRum*) as a flexible multi-channel access mechanism, and employs P-CSMA (*Probabilistic CSMA*) to temporally adapt the use of CSMA to the intensity of collision.

The major principle behind FAVOR is a paradigm shift from *discrete channel allocation* to *continuous frequency allocation*. Specifically, we treat the spectrum resource as a continuous domain rather than discrete channels, based on which FAVOR transforms the *frequency allocation* problem into a *spatial tessellation* problem by unifying frequency and distance into one continuous space, and adapts CVT (*Centroidal Voronoi Tessellation*) method [7] to seek a near-optimal solution that the nodes/links that are closer to each other in distance are further away from each other in frequency. Moreover, the density function used in CVT enables us to take as input the intensity of WiFi interference. Though FAVOR can be viewed as a continuous version of the conventional graph coloring approach, it offers a much greater freedom due to the relaxation from a discrete set to a continuous space.

Although FAVOR's flexible spectrum access enables ZigBee links to fit the "frequency white space" of WiFi spectrum, interference may not be totally avoidable. This is where P-CSMA works to opportunistically leverage the "temporal white space" of the time-varying interference. Based on real-time PRR measurements fed by a receiver to indicate the interference intensity, the randomized control of P-CSMA gives a better chance to disabling CSMA under a good link quality whereas it gets biased towards enabling CSMA under a severe interference. In a nutshell, P-CSMA enables ART to fine-tune the tradeoff between throughput and PRR under a time-varying interference that cannot be totally avoided by FAVOR.

In summary, we make the following main contributions in this paper as follows.

- We propose to shift the paradigm from traditional discrete channel allocation to continuous frequency allocation.
- We innovate in transforming frequency allocation into spatial tessellation in a frequency-distance space, and propose an algorithm to compute a near-optimal solution in a localized manner.
- We design a probabilistic mechanism to adaptively use CSMA according to real-time interference assessments; it can fine-tune the tradeoff between transmission throughput and quality.
- We perform extensive experiments in order to demonstrate the efficacy of our proposal in tackling the ZigBee-WiFi co-existence issue.

The remainder of our paper is organized as follows. We first report the experiments motivating our design of ART in Sec. 2. We then present our ART framework in detail in Sec. 3. The evaluation of ART is given in Sec. 4, and the most related literature is surveyed in Sec. 5. We finally conclude our paper in Sec. 6.

2 Motivations and Background

In this section, we first report the tradeoff between frequency and distance in Sec. 2.1, to motivate the innovation of FAVOR in flexible spectrum access. We then illustrate the impact of WiFi interference on ZigBee spectrum and hint our later solution to handling WiFi interference in Sec. 2.2 and Sec. 2.3, respectively. We finally introduce the mathematical background on which our FAVOR algorithm will be based in Sec. 2.4.

All the experiments presented in this section are performed using Micaz motes, and evaluated by the following two metrics.

- **Throughput**. In this paper, we measure the throughput of a transmission link as the ratio between the receiving data rate (i.e., the number of bits received in one time unit) at the receiver and the nominal data rate at the transmitter. For example, in our experiments, the nominal data rate of MicaZ mote is 250 kbps.
- **Packet Reception Ratio** (or PRR). We hereby adopt standard definition of PRR: the radio between the number of received data packets and the number of transmitted data packets.

2.1 Frequency-Distance Tradeoff

To demonstrate the feasibility of treating frequency spectrum in a continuous manner, we set up an experiment shown in Fig. 2: two parallel ZigBee links l_1 and l_2 are set by two MicaZ motes. We fix the location of l_1, and change the one of l_2 such that the distance between l_1 and l_2 varies from 1.2 m to 4.8 m. We let the transmitters send data packets as fast as possible with identical transmission power. As we are interested in ZigBee interference for now, we operate the links from 2475 to 2480 MHz to avoid WiFi interference.

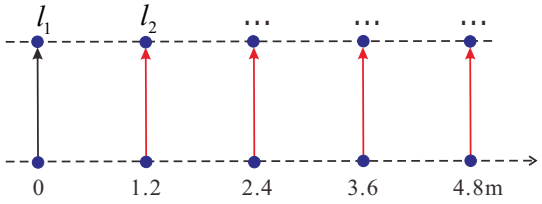

Fig. 2. Experiment configuration. The distance from transmitter to receiver is 3.6m.

We hereby fix the frequency of l_1 at 2480 MHz, while changing the frequency of l_2 from 2475 to 2480 MHz. The throughputs and PRRs for all the frequency-distance combinations are shown in Fig. 3 (a) and (b), respectively. The figures clearly show throughput and PRR tradeoff surfaces that we may achieve by extending from a 1D discrete frequency space (for conventional channel allocations) to a higher dimensional *frequency-distance space*.

We also plot two sets of cut views of Fig. 3 in Fig. 4. In Fig. 4 (a) and (b), the distance is fixed at 1.2 m, and

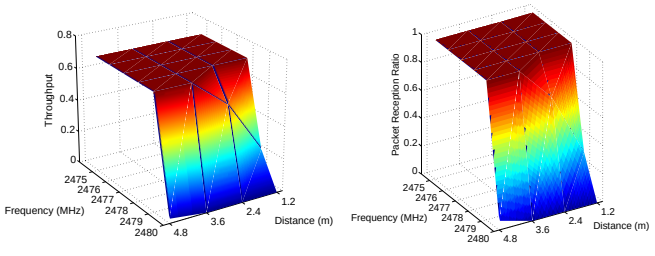

(a) Throughput (b) PRR

Fig. 3. The throughput and PRR for the worse link with different frequency-distance combinations.

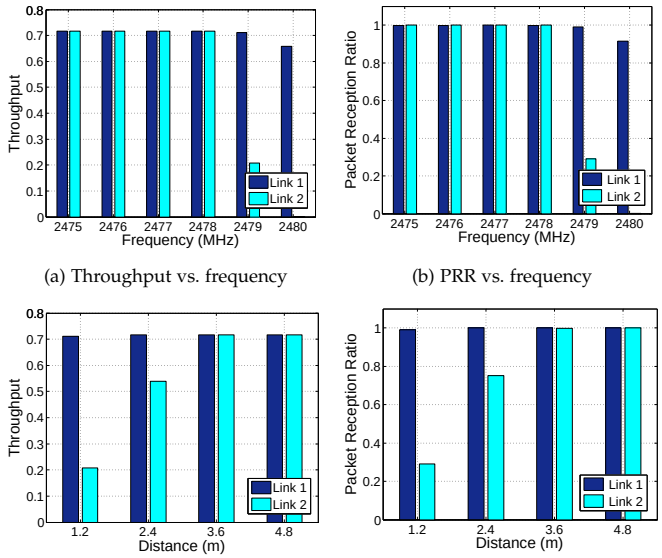

(a) Throughput vs. frequency (b) PRR vs. frequency

(c) Throughput vs. distance (d) PRR vs. distance

Fig. 4. The throughputs and PRRs of both links as functions of frequency (a), (b) and distance (c), (d).

in Fig. 4 (c) and (d), the frequency of l_2 is fixed at 2479 MHz. As shown in Fig. 4 (a) and (b), when two consecutive frequencies (e.g., 2479 and 2480 MHz) are used, the resulting channel overlapping ruins one of the links: one "captures" the transmission medium thereby achieving a nearly full throughput (beyond 0.7) and a 100% PRR, while the other one gets much lower throughput. Under the same short distance, a frequency separation of 2 MHz (e.g., 2478 and 2480 MHz) is usable for achieving a superior throughput and PRR. Fig. 4 (c) and (d) further demonstrate, by increasing the distance between two links (e.g., 3.6m), even a frequency difference of 1 MHz is sufficient to achieve an almost full throughput and a perfect PRR. In other words, slightly increasing distance can resolve the interference resulting from channel overlapping.

One may suggest to employ CSMA to resolve the interference between overlapped channels. But the overhead of CSMA heavily restricts the throughput, e.g., if we repeat the above experiments on the two links both in 2480 MHz with CSMA enabled to resolve the resulting co-channel ZigBee interference, the throughput is decreased to 0.16 for both of them, while our proposal of frequency-distance tradeoff can achieve a 4.5 times higher throughput and the same good PRR. According to our experience, the temporal overhead of CSMA (and thus the induced energy consumption) may

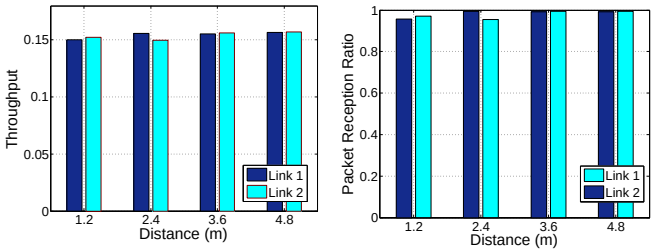

(a) Throughput (b) PRR

Fig. 5. The throughputs and PRRs for two CSMA-enabled links with different frequency-distances combinations.

drastically threaten sensor systems with extremely limited energy supply (e.g., [22]).

Recall that ZigBee standard [2] support only two orthogonal channels in the spectrum 2475-2480 MHz, whereas we almost get six overlapped but highly usable channels by varying center frequency and link distance in a continuous manner.[1] Therefore, another important message we get is:

> *Continuously allocating frequencies to links with respect to the distances between them allows for a much more flexible and efficient use of limited spectrum, which brings great potential to improve network throughput and to guarantee transmission quality.*

The above observations is essential to the design of FAVOR (see Sec. 3.2). Now the question is how the performance is affected by co-existing WiFi networks, which we will further investigate in Sec. 2.2.

2.2 WiFi Interference to ZigBee under Continuous Spectrum

To reveal the performance of CSMA in face of WiFi interference, we perform our experiments on a single ZigBee link in our research center where many WiFi networks exist (Fig. 1). We continuously tune its center frequency to scan the whole WiFi spectrum (i.e., 2405-2474 MHz) instead of only the 16 standard ZigBee channels. To verify the temporal variation of WiFi interference, we measure throughput and PRR every one hour during the 8 working hours in a day. We disable CSMA to illustrate the impact of WiFi interference to ZigBee transmissions across the spectrum, and enable CSMA to demonstrate the coordination effect of CSMA under WiFi interference. Fig. 6 (a) and (b) shows the case with CSMA disabled. Since the transmitter sents data packets without carrier sensing at a constant transmission rate across the spectrum, a lower PRR (or throughput) implies heavier WiFi interference. It is clearly demonstrated, the WiFi interference mainly comes from three orthogonal 802.11 channels (i.e., channels 1, 6 and 11). In the worst case, the PRR can go down below 0.5. In contrast, in the frequencies where WiFi interference is weak or absent (e.g., around 2425, 2450 and 2474 MHz), the throughput is rather high and the PRR is almost perfect. In Fig. 6 (c) and (d), we turn on CSMA to deal with WiFi interference. As shown in Fig. 6 (c), the throughput is quite low across the spectrum, even in those frequencies less interfered by WiFi, which further confirms

1. If the center frequency of CC2420 radio can be tuned at a finer granularity (e.g., of less than 1 MHz) in the future, our mechanism will offer more available channels.

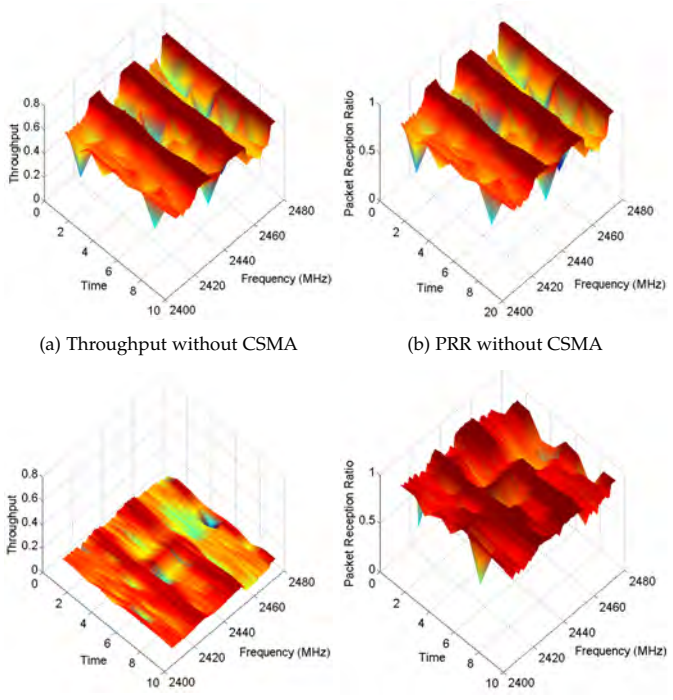

(a) Throughput without CSMA (b) PRR without CSMA

(c) Throughput with CSMA (d) PRR with CSMA

Fig. 6. The impact of CSMA on throughput and PRR in different frequencies under WiFi interference.

the restriction of CSMA on transmission data rate we have illustrated in Sec. 2.1. Referring to Fig. 6 (d), we find that a rather high PRR is offered due to the effective coordination provided by CSMA even in those heavily interfered frequencies (e.g., around the frequencies of 2412, 2437 and 2474 MHz). Although the price in throughput we pay for PRR may be reasonable in those frequencies where there is severe WiFi interference, it is a waste for the less or non interfered frequencies where we have the possibility of gaining high throughput.

Another observation is that, though the intensity of WiFi interference has small-scale variation for a short moment, it is (relatively) stable for a long run, according to the regular shapes of the waveform surfaces shown in Fig. 6 (a) and (b). Therefore, it suggests a FDMA-CSMA hybrid strategy. Recalling that our experiment results presented in Sec. 2.1 have suggested an approach to better exploit the spectrum of high quality averagely (the one corresponding to the peak of the waveform surface), we hereby have the following insight in terms of utilizing CSMA to tackle the short-term variation of WiFi interference.

> *Blindly enabling CSMA across the whole spectrum does not adapts well to WiFi interference. The use of CSMA should be tuned in a more flexible fashion according to the WiF interference in a specific frequency band.*

2.3 Adapting to the Temporal White Space of WiFi Interference

Inspired by the above observations, we study the possibility of a randomized control on disabling/enabling CSMA. The basic idea is, for every packet to be sent, a coin is tossed based on a certain probability (namely *CSMA Probability*) to decide whether to enable CSMA or not. In this experiment, we operate a single ZigBee link under two interfered frequencies (2455 and 2460 MHz) and a clean frequency (2480 MHz).

Fig. 7(a) shows that for all the three frequencies, the throughput drops with an increasing CSMA probability, consistent with our observation in Sec. 2.1. The throughputs

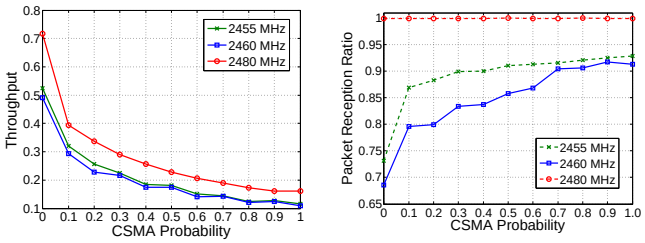

(a) Throughput (b) PRR

Fig. 7. Throughputs and PRRs under different CSMA probabilities.

for 2455 and 2460 MHz are lower than the one for 2480 MHz due to the ambient WiFi interference. An opposite trend can be observed in Fig. 7(b), where increasing the CSMA probability may significantly improve PRR for the two interfered frequencies. Our results in Fig. 7 demonstrate that a randomized control does allow for a smooth fine-tuning of the tradeoff between throughput and PRR. For example, in the 2455 MHz channel, a CSMA probability of 0.3 doubles the throughput with respect to the totally enabled CSMA (i.e., CSMA probability is 1), while achieving a PRR of 0.9 (very close to the totally enabled CSMA). If we afford a PRR of 0.85, we can take a CSMA probability of 0.1 and hence increase the throughput by a factor of three compared with the totally enabled CSMA. So the main insight is:

> *Adding a randomized control to determine whether to enable CSMA or not on a per packet basis endows us with the leverage to fine-tune the tradeoff between throughput and PRR.*

2.4 Mathematical Background on Spatial Tessellation and CVT

As a theoretical basis of our later algorithm, we hereby briefly introduce a few key ideas on spatial tessellation. Given a region $\mathcal{A} \subseteq \mathbb{R}^n$, the set $\{A_i\}$ is called a *tessellation* of $\mathcal{A}$, if $A_i \cap A_j = \emptyset$ for $i \neq j$ and $\cup_i A_i = \mathcal{A}$. Let $\|\cdot\|_{\ell^2}$ denote the Euclidean norm on $\mathbb{R}^n$. For a set of points $\{u_i\}$ belonging to $\mathbb{R}^n$, the Voronoi region V_i corresponding to the point u_i is defined by

$$V_i = \left\{v \in \mathcal{A} | \|v - u_i\|_{\ell^2}^2 \leq \|v - u_j\|_{\ell^2}^2, \forall j \neq i\right\}.$$

The set $\{V_i\}$ is termed *Voronoi tessellation* of $\mathcal{A}$, with points $\{u_i\}$ called *generators* and each V_i referred to as the *Voronoi cell* corresponding to u_i. An energy function is defined upon each generator according to its Voronoi cell. One typical metric of measuring such an energy is $\|v - u_i\|_{\ell^2}$ where $v \in V_i$, hence generator u_i has an energy $\int_{V_i} \|v - u_i\|_{\ell^2}^2 \Phi(v) dv$ where $\Phi(v)$ denotes the density function at v. It is obvious that the optimal location for node n_i to minimize its energy

is the centroid of V_i. Moreover, a generator can impact other nearby generators. In particular, if a generator has high energy, i.e., deviating from the centroid and getting close to a boundary edge of its Voronoi cell, it exerts an impact to another one whose Voronoi cell shares this boundary edge. Such mutual impact can be characterized by the dual graph of the Voronoi tessellation, formed by using edges to connect the pairs of the generators whose associated Voronoi cells are adjacent. The generators connected by shorter edges in the dual graph have larger mutual impact.

The aim of *Centroidal Voronoi Tessellation* (CVT) is to calculate an optimal tessellation by varying the locations of the generators (and thus their Voronoi tessellation) such that their total energy (1) is minimized .

$$\Im(\{V_i\},\{u_i\}) \quad = \quad \sum_i \int_{V_i} \|v-u_i\|_{\ell^2}^2 \Phi(v)dv. \qquad (1)$$

As this objective is not convex, achieving a global optimality is hard. Fortunately, the CVT theory [7] indicates a good locally optimal solution that every u_i coincides with the centroid of its Voronoi cell V_i. Consequently, the generators in CVT are evenly spread in the space if the density function is constant and have almost the same energy. In other words, CVT leads to an energy-balanced tessellation. From the dual point of view, minimizing the total energy of the generators is equivalent to minimizing their mutual impact. Moreover, the distribution of the generators well adapts to a non-constant density function $\Phi(v)$: a larger function value leads to a higher generator density. These unique properties of CVT are related to our need for "spreading" sensor nodes (or links) over the frequency spectrum with their "mutual impacts" (i.e., interference) minimized. In fact, our FAVOR algorithm is a non-trivial extension of CVT to a space involving both frequency and distance (or location).

One classic approach of calculating CVT is Lloyd's method [15], which iteratively computes the Voronoi tessellation and moves the generators to the centroids of their associated Voronoi cells. However, when the input is in a high-dimensional space (e.g., $\mathbb{R}^n$ for $n > 2$), computing Voronoi tessellation turns out to be highly non-trivial. A very complicated data structure is required to model a $\mathbb{R}^3$ Voronoi tessellation, and the induced computational cost is too high to be affordable for hardware platforms with limited energy and computing capability (e.g., MicaZ motes). For $\mathbb{R}^n$ $(n > 3)$, no algorithm for CVT has been implemented even for common CPUs.

3 ART: Co-Existing of ZigBee and WiFi in Frequency and Time

In this section, we first discuss our system model. Then we present in detail our FAVOR algorithm and P-CSMA protocol, respectively. We finally show how ART can be implemented in a practice scenario.

3.1 System Model

We consider a WSN consisting of a set of ZigBee nodes $\mathcal{N} = \{n_1, n_2, ..., n_N\}$ with $|\mathcal{N}| = N$, which are deployed on a 2D plane. Although our proposal can be extended to a 3D space, we confine our scenarios to 2D due to the limitation of our experiment conditions. Let $\{u_i\} \subseteq \mathbb{R}^2$ denote the locations of the nodes. We denote by $\mathcal{N}(n_i)$ the one-hop neighbors of node n_i. Node n_i can utilize range information of its one-hop neighbors to construct a local coordinate system [17], if there is no positioning device equipped.

Given a frequency band $\mathcal{B} = [f_{min}, f_{max}]$ and the channel width f_w, we can assign each nodes n_i a center frequency $f_i \in \mathcal{B}'$, where $\mathcal{B}' = [f_{min} + f_w/2,\ f_{max} - f_w/2]$. Moreover, we introduce a density function $\Phi(f)$ (with $f \in \mathcal{B}'$) which defines the qualities of the frequencies. In particular, the frequencies less interfered by WiFi have higher density values. This density function can be updated and disseminated by a central station (e.g., a sink node) within the network. Since the quality for a frequency is stable for a long run (as shown in Sec. 2.2), we do not need to update it very often unless there is a drastic change in the surrounding WiFi deployment. Considering both location and frequency, the representations of nodes are extended from 2D to 3D space, i.e., each node has a new "coordinate" $(u_i, f_i) \in \mathbb{R}^2 \times \mathcal{B}'$.

Each node n_i maintains a *CSMA probability* $Pr_{csma}(n_i) \in [0, 1]$ indicating the probability of sending a packet with CSMA enabled. We denote by β the granularity of tuning $Pr_{csma}(n_i)$. We also use *qualification range* in PRR: $\mathcal{P} = [P_{min}, P_{max}]$ to describe the performance required by specific applications. The data transmission is said to be *qualified* if the measured PRR prr falls into the range, i.e., $prr \in \mathcal{P}$.

3.2 FAVOR: A Location-Aware Frequency Allocation Scheme

We hereby present our FAVOR algorithm, including its optimization objective and the algorithm to find a locally optimal solution. We also discuss a practical implementation issue at the end.

3.2.1 Objective: Balancing Distance and Frequency

According to our observations in Sec. 2.1, a frequency allocation scheme should assign very different frequencies to two nodes that are close to each other but arbitrary frequencies to nodes that are far from each other. A possible location-dependent frequency allocation is illustrated in Fig. 8.

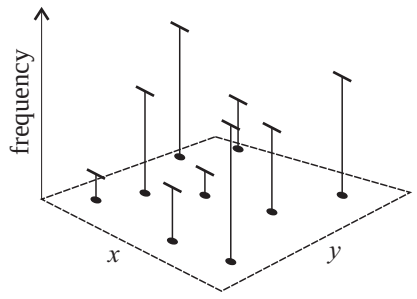

Fig. 8. Allocating (center) frequencies to a 2D WSN based on the node location. The black points on the 2D plane denotes the nodes, while the "poles" of different heights on the points indicate different frequencies.

We define the following *energy metric* for each sensor node based on Euclidean norm in the frequency-distance space: $\delta_i\|v-u_i\|_{\ell^2} + \|f-f_i\|_{\ell^2}$, where $v \in \mathbb{R}^2$ and $f \in \mathcal{B}'$ and δ_i is a factor to characterize the interference effect under certain transmit power for node n_i. We divide the location-frequency space $\mathcal{A} = \mathbb{R}^2 \times \mathcal{B}'$ into Voronoi tessellation $\{V_i\}$

with $n_i = (u_i, f_i)$ serving as generators. Each node n_i has an energy

$$F_i(V_i, f_i) = \int_{V_i} \left(\delta_i \|v - u_i\|_{\ell^2}^2 + \|f - f_i\|_{\ell^2}^2\right) \Phi(f) dz \quad (2)$$

where $z = (u, f) \in V_i$. Since the sensor nodes usually take the maximum transmit power to resist high-power WiFi interference, we assume $\delta_i = 1$ in our case. The interference relationship can be represented by the dual graph of the Voronoi Tessellation: nodes close in both frequency and distance have adjacent Voronoi cells and are thus connected in the dual graph, which implies they may interfere each others.

We aim at an optimal "tessellation" of $\mathcal{A} \in \mathbb{R}^3$ such that the following FAVOR objective is minimized.

$$\mathfrak{F}(\{V_i\}, \{f_i\}) = \sum_i F_i(V_i, f_i). \quad (3)$$

As mentioned in Sec. 2.4, this is equivalent to minimizing the interference among the sensor nodes according to the dual graph. It is obvious that (3) differs from (1) mainly in that part of the coordinates are fixed: we do not get the freedom to move nodes, only the frequencies allocated to them are variables. Fortunately, the absence of the freedom does not compromise the ability of CVT in minimizing interference. According to what we discussed in Sec. 2.4, an optimal solution tends to spread out the sensor nodes in the frequency-location space. Since the locations in $\mathbb{R}^2$ are fixed, minimizing FAVOR objective function can be done only by separating the closed nodes in the frequency dimension. In particular, frequencies for two close-by nodes will be pushed further from each other. Moreover, through the density function, we can encourage the nodes to choose the frequencies less interfered by WiFi; thereby reducing the reliance on CSMA to deal with WiFi interference. Additionally, Equation (3) is meant for only one contention domain. When taking into account the existence of obstacles which separate the sensor network into multiple contention domains (e.g. the indoor WSNs may be deployed in different rooms, i.e., contention domains, such that the sensor nodes are separated by walls, as we will shown in Sec. 4.4), different objective functions can be defined in each contention domain where our algorithm can be individually performed.

3.2.2 Algorithm: A CVT-based Approach

The challenge we face now is twofold: i) as we need to perform CVT at least in a 3D space (it can be 4D if nodes are distributed in a 3D space), we need an algorithm more efficient than the Lloyd's method [15]; otherwise, nodes with limited computation resource cannot afford it, and ii) part of the coordinates for each $u_i \in \mathcal{A}$ are fixed, whereas CVT requires $\{u_i\}$ to be variables. Hence, we first propose *Approximate CVT* (A-CVT) to transform the problem into a more tractable and implementable form, then apply gradient projection method to handle the fixed coordinates.

Given a region $\mathcal{A} = \mathbb{R}^2 \times \mathcal{B}' \subseteq \mathbb{R}^3$ and suppose we apply Voronoi tessellations to partition $\mathcal{A}$, we may re-write the objective (3) as the following.

$$\mathfrak{F} = \int_A \min_i \left(\|v - u_i\|_{\ell^2}^2 + \|f - f_i\|_{\ell^2}^2\right) \Phi(f) dz \quad (4)$$

The equivalence between (3) and (4) is obvious: for each generator (u_i, f_i), integrating over its own cell V_i implies an integration over all the points in $\mathcal{A}$ that are closer to (u_i, f_i) than to any other generators (by the definition of a Voronoi cell in Sec. 2.4). Now we get a global integration over $\mathcal{A}$, eliminating the need for re-computing the Voronoi tessellations in every iteration. Furthermore, we apply an approximation to $\min(\cdot)$ to make the problem tractable, leading to the following A-CVT objective.

$$\mathfrak{F}(\{V_i\}, \{f_i\}) = \int_A \left[\sum_i \left(\left(\|v - u_i\|_{\ell^2}^2 + \|f - f_i\|_{\ell^2}^2\right) \Phi(f)\right)^\lambda\right]^{\frac{1}{\lambda}} dz. \quad (5)$$

The proof of (5) converging to (4) when $\lambda \to -\infty$ is rather trivial and is thus omitted. In practice, we take $\lambda \in [-40, -20]$.

As minimizing the A-CVT objective (5) is a typical nonlinear optimization problem, we apply a gradient-descent method with gradient projection to search for a local minimum. The pseudocodes of the algorithm are shown by **Algorithm 1**. Each sensor node n_i first measures the distances to the nodes belonging to $\mathcal{N}(n_i)$, based on which a certain space embedding technique (e.g. [17]) is applied to construct a local coordinate system (line 1). This step is not necessary if additionl positioning mechanisms/devices are available. Roughly, the main part of the algorithm proceeds in rounds

Algorithm 1 FAVOR

Input: For each $n_i \in \mathcal{N}$, initial frequency $f_i^0 \in \mathcal{B}'$, stopping tolerance ε_1 and ε_2
Output: f_i^* for each n_i
1: Construct a local coordinate system to obtain the locations of $\mathcal{N}(n_i)$
2: For every node $n_i \in \mathcal{N}$ in each round (every τ ms):
3: Compute the gradient $g(z_i)$ of (5)
4: Project $g(z_i)$ on f_i to get $g_f(z_i)$
5: $f_i^+ = f_i - \alpha \cdot g_f(z_i)$ /*α is the step size*/
6: **if** $|g_f(z_i)| > \varepsilon_1 \vee |f_i^+ - f_i| > \varepsilon_2$ **then**
7: $f_i \leftarrow f_i^+$; BROADCAST(f_i) to nodes in $\mathcal{N}(n_i)$
8: **else**
9: $f_i^* \leftarrow f_i^+$
10: **end if**

and takes the following three steps in each round:

1) Compute gradient for each generator $z_i = (u_i, f_i)$ as (line 3)

$$g(z_i) = \int_{\mathcal{A}} 2\left(\Phi(f)\|z - z_i\|_{\ell^2}^2\right)^{\lambda-1} \Phi(f)(z_i - z) \left(\sum_j \left(\Phi(f)\|z - z_j\|_{\ell^2}^2\right)^\lambda\right)^{\frac{1-\lambda}{\lambda}} dz \quad . \quad (6)$$

In order to facilitate localized computation and also to reduce the complexity, the summation in the third term can be applied only for $j : n_j \in \mathcal{N}(n_i)$ and the integration can be done only for v in the neighborhood of z_i. This is possible because the terms introduced by those far-away locations contribute only insignificantly to $g(z_i)$, due to the very small

value of λ. This is also intuitively correct as the change in z_i for CVT is only affected by z_j whose cell shares boundaries with that of z_i. For a network, we can use the communication neighborhood to approximate the tessellation neighborhood.

2) As u_i is fixed and only f_i is variable, we take $g_f(z_i)$ as the projection of $g(z_i)$ on the frequency axis (line 4).
3) A tentative update is applied to the frequency by $f_i^+ = f_i - \alpha \cdot g_f(z_i)$ where α is a step size. If both $|g_f(z_i)|$ and $|f_i^+ - f_i|$ become sufficiently small, the algorithm is terminated, returning the optimal frequency allocation (line 9); otherwise, the frequency of each n_i is updated by $f_i \leftarrow f_i^+$, the outcome is exchanged among neighboring nodes (line 7), and further compuation will be conducted during the next round.

We omit the convergence analysis as it follows directly from the basic theory of gradient-descent methods [3]. We also note that, if a centralized computing is allowed, we can use the Quasi-Newton method to solve (5), which often leads to a faster convergence.

We take a one-dimensional network[2] as an example and illustrate the result of FAVOR in Fig. 9. To show the effect of the density function $\Phi(f)$, we assume a density function $\Phi(f)$ increasing in f (indicated by the color spectrum, i.e., higher frequencies marked by warm color are less interfered by WiFi while lower frequencies marked by cold color suffer from more WiFi interference). All nodes are initially assigned ascending frequencies according to the node IDs with small frequency separations. The outcome of FAVOR shows that the neighboring nodes are assigned frequencies that are significantly further apart (especially for more interfered frequencies), while more nodes choose higher frequencies to avoid being interfered by WiFi.

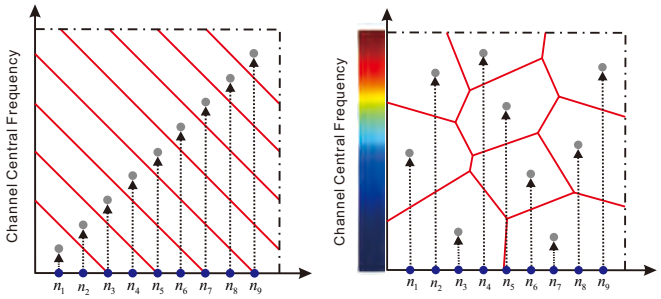

(a) Initial frequency allocation (b) Outcome of FAVOR

Fig. 9. Frequency allocation based on Voronoi tessellations. A density function $\Phi(f)$ is defined on the frequency-axis. The warm and cold colors indicate the high and low values of $\Phi(f)$, respectively. The Voronoi tessellations are also shown by red line segments.

3.2.3 Extension: Frequency Allocation for Links

As FAVOR relies on a set of (point) locations $\{u_i\}$, an obvious difficulty it may face is what if no obvious points exist in a networking scenario. In particular, a network usually consists of several point-to-point wireless links. We propose two possible solutions for centralized and distributed computing respectively. If a centralized computing is feasible, we may apply the extended Voronoi diagram where generators are not points but line segments (representing the links) [4]. This method may result in rather accurate frequency allocation, but it is quite time consuming. Therefore, we pick one point to represent each link in a distributed computing environment. This point can be either the source or the destination of the link, or it can even be the middle point of the link. In the former case, the computation is performed by that node, while in the latter one, the computation can be done by either of the two nodes. Our experiments reported in Sec. 4.3 show that the distributed (point) approach still leads to very good performance.

3.3 P-CSMA: Randomizing the use of CSMA

Although FAVOR allows us to better utilize the spectrum resource, exploiting only the "clean" spectrum (2475 to 2480 MHz) may not be sufficient for densely deployed ZigBee networks; hence, certain ZigBee links have to step into the WiFi interfered frequency band, e.g., those frequencies in the non-peak areas shown by Fig. 6. Therefore, CSMA is still necessary to guarantee transmission quality, but it should be deployed in a more flexible manner. Motivated by the observations presented in Sec. 2.2 and 2.3, we propose P-CSMA to achieve a balanced performance between throughput and PRR through a randomized control.

In P-CSMA, each transmitting node enables CSMA to send a packet with probability Pr_{csma}. Specifically, Pr_{csma} is tuned up to maintain transmission quality at a minor cost of throughput when experiencing severe WiFi interference; otherwise it is tuned down to improve throughput. It has been reported that WiFi interference exhibits temporal variation [10], but according to our experience, the temporal variation of WiFi interference is highly unpredictable, especially when multiple WiFi networks exist. Therefore, to identify the temporal white space of WiFi interference, we opt for a real-time evaluation method in our P-CSMA protocol, rather than the model-driven predication method proposed in [10], such that a transmitter adaptively tunes Pr_{csma} based on the real-time interference assessment provided by the receiver. In the following, we explain in detail how the real-time assessment is conducted and how Pr_{csma} is tuned accordingly.

3.3.1 Real-Time Interference Assessment

As mentioned earlier, we use PRR to indicate the interference level. A PRR is computed by a receiving node and is reported back to the transmitting node periodically. We disable ACK for most of the packets but only enable it for the receiving node to send a PRR report. Data packets delivered by P-CSMA are grouped into *transmission windows* of size W, i.e., each window contains W data packets, and each data packet carries an ID field to indicate its source and the transmission window to which it belongs. We denote by $pkt^i_{j,k}$ the j-th data packets involved in the k-th transmission window of node n_i, where $k = 1, 2, \cdots$ and $j = 1, 2, \cdots, W$. Data packets not belonging to a window have $k = 0$. An example of P-CSMA is shown in Fig. 10, where n_t and n_r denote the transmitting and receiving nodes, respectively.

2. Our algorithm actually applies to 2D and 3D cases, but we only show the results of a 1D deployment to avoid visual confusion.

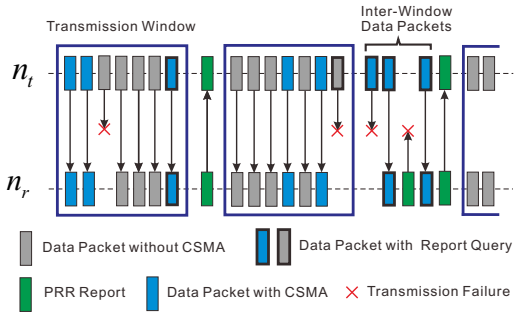

Fig. 10. An example of P-CSMA. In each transmission window, whether CSMA is enabled for individual data packets is determined by the current CSMA probability. If n_r correctly receives the last data packet in a transmission window (a DPRQ), it will feed the PRR of this window back to n_t. In case that this DPRQ or its corresponding PRR report gets loss, n_t keeps sending inter-window data packets (another type of DPRQ) with CSMA deterministically enabled until the expected PRR report arrives. For every DPRQ from n_t, n_r responds with the PRR of the latest window.

A transmitting node n_t sends the data packets in a transmission windows with its current CSMA probability $Pr_{csma}(n_t)$. Upon sending out all packets in a window, n_t expects a PRR report piggyback with an ACK from the receiving node n_r, and then updates $Pr_{csma}(n_t)$ accordingly. When waiting for the PRR report, n_t does not suspend its packet transmissions; it instead keeps transmitting *inter-window data packets* (i.e., the packets between two successive transmission windows) but with CSMA deterministically enabled, so as to ensure the execution of PRR reporting on one hand and to avoid interfering the transmissions of PRR reports in opposite direction on the other hand. Finally n_t starts the next transmission window with the newly updated $Pr_{csma}(n_t)$.

By checking the ID fields of the data packets, the receiving node n_r can identify the last data packet in the current transmission window, $pkt^t_{W,k}$ ($k \neq 0$). Upon receiving this tail packet, n_r feeds a PRR report back to n_t by attaching it to an ACK. The PRR for each transmission window can be computed by dividing the number of received packets by W. To ensure a successful PRR reporting, n_r reports PRR upon every inter-window packet (if any) that can be identified by checking if $k = 0$, in case that the tail packet in the past window or the previous PRR reports get lost. These data packets (including $pkt^t_{W,k\neq 0}$ and $pkt^t_{j,0}$) that can trigger n_t's PRR reporting behavior are termed *Data Packets with Report Query* (DPRQ) in Fig. 10.

3.3.2 Tuning CSMA probability

Being aware of the immediate PRR, the transmitting node n_t adapts its CSMA probability $Pr_{csma}(n_t)$ accordingly, and then starts a new transmission window. The objective of updating $Pr_{csma}(n_t)$ is to maintain the PRR of next transmission window in the predefined qualification range $\mathcal{P}$. In our implementation, $Pr_{csma}(n_t)$ is tuned in an additive manner with a step size β in Equation (7).

$$Pr^{+}_{csma} = \begin{cases} \min\{Pr_{csma} + \beta, 1\} & \text{if } prr < P_{min} \\ \max\{Pr_{csma} - \beta, 0\} & \text{if } prr > P_{max} \\ Pr_{csma} & \text{otherwise} \end{cases} \quad (7)$$

Specifically, if the measured PRR is below our predefined qualification range ($prr < P_{min}$), we increase $Pr_{csma}(n_t)$ to guarantee the transmission quality in face of severe WiFi interference. Otherwise, if the PRR is beyond the qualification range ($prr > P_{max}$), the current $Pr_{csma}(n_t)$ is too conservative. So we decrease its value to gain a higher throughput. According to our earlier experiment results (shown in Fig. 7), both throughput and PRR are very sensitive to $Pr_{csma}(n_t)$ in the range of $[0, 0.1]$. To allow for a finer granularity in adjusting $Pr_{csma}(n_t)$, our implementation takes $\beta = 0.01$ if $Pr_{csma}(n_t) \in [0, 0.1]$; otherwise, $\beta = 0.1$.

3.4 Application: Tree-based Data Collection

Data collection is a fundamental operation in WSNs. During a data collection process, a set of nodes $\{n_i\}$ send their collected data towards a common sink node n_s, usually over a tree-based routing topology with n_s as the root (e.g., [8]). Such a tree-based data collection strategy has been recently integrated into applications of mobile sensing. For example, in [11], the data sensed by mobile leaf nodes are carried to users through a tree-based backbone wireless network consisting of ZigBee nodes.

When applying FAVOR as a multi-channel access mechanism, we need to handle the conflicts incurred by multiple links sharing the same incident node, as such conflicts are not the consequence of co-channel interference but of a node having only one radio. Therefore, we need a TDMA-like time schedule to assist ART with resolving the type of conflicts. While a joint frequency allocation and scheduling problem is beyond the scope of our paper, we simply adapt the minimum latency scheduling mechanism proposed in [20]. We basically arrange the nodes into layers according to their hop-distances to the sink on the collection tree T, then we use a labeling mechanism similar to [20]. The idea of this labeling is twofold: i) to guarantee that the number of labels assigned to an outgoing link of node n_i should be 1 plus the number of descendants of n_i in T, and ii) to assign different labels for links sharing the same node. Whereas the proposal in [20] adopts (orthogonal) channel allocation to enable parallel transmissions and allocates channels using a first-fit distance-$(\rho+1)$ graph coloring, our FAVOR allocates overlapped channels and thus potentially allows for more parallel transmissions.

Under ART, each node except leaf nodes is assigned with a center frequency by FAVOR for receiving the data packets from its child nodes, whereas the leaf nodes do not need their own receiving frequencies. At a given point in time, for those links that are activated according to the time schedule discussed above, both receivers and transmitters tune to the receiving frequencies assigned to the receivers. Every pair of nodes forming an active link perform P-CSMA during their data transmissions so as to adapt the use of CSMA to WiFi interference. Although our ART framework can be adapted to other communication patterns, e.g., broadcast and aggregation, we focus only on data collection in this paper.

4 Evaluation

We have implemented ART in our MicaZ platforms and performed extensive experiments. We first present the basic configuration for our experiments, and then report the

results in different network scenarios. Finally, we briefly examine the convergence of our distributed FAVOR algorithm.

4.1 Experiment Settings

We adopt MicaZ motes and TinyOS 2.1 as the hardware and software platforms, respectively. We perform our experiments within the frequency spectrum of $\mathcal{B} = [2449, 2459]$ MHz that is partially interfered by the WiFi networks already deployed in our research center. As we will shown in our experiments, such a narrow spectrum is sufficient to support highly parallized transmissions in our small-scale but dense WSN. As a ZigBee channel has a width of 2 MHz, the center frequencies we can employ are actually within the ranges of $\mathcal{B}' = [2450, 2458]$ MHz. According to the ZigBee standard [2], there are only two channels specified within this spectrum, i.e., Channel 20 and 21 with center frequencies at 2450 and 2455 MHz, respectively. We term this channel allocation scheme *two-channel*. We also test the proposal in [25], where a frequency separation of 3 MHz is suggested. Given the available spectrum, we may have at most three channels (e.g., those with center frequencies at 2450, 2453 and 2456 MHz), and we term this scheme *three-channel*. We apply a greedy graph coloring approach to allocate channels for both two- and three-channel schemes, while our ART employs FAVOR to allocate all nine potential channels (center frequencies) to nodes. In the following, we use only center frequency to indicate a channel. Considering PRR as a metric to measure frequency quality, we define a density function $\Phi(f)$ for our FAVOR algorithm based on the results shown in Fig. 6(d). which ends up with the density function shown in Fig. 11.

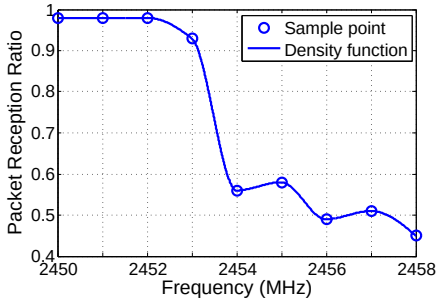

Fig. 11. Density function $\Phi(f)$ ($f \in [2450, 2458]$ MHz).

We deploy our WSNs on the ceiling of our laboratory (see Fig. 12), emulating an indoor monitoring application scenario. We perform each experiment for 5 minutes, and use the *Total Received Packets* (TRP) and PRR at each destination node as the performance measures. Note that, the value of TRP is actually an indicator of throughput, while PRR is the metric for transmission quality. We set the qualification range of PRR as $\mathcal{P} = [0.85, 0.9]$, and initialize the CSMA probabilities to 0.2 for all nodes.

4.2 ART on A Single Link

We first verify the efficacy of ART on a single link. As we do not need channel allocation in this case, only P-CSMA is applied on this link. We let the transmitting node persistently send data packets to the receiver as fast as possible at the center frequency 2455 MHz, which suffers from rather severe WiFi interference (as shown in Fig. 11). Besides, we vary the transmission window size ω from 100 to 500 to investigate how ω impact the performance of ART. We also compare our ART with *BuzzBuzz* [14], a MAC layer protocol for ZigBee to survive WiFi interference, which can be performed by individual links with CSMA enabled. As a benchmark, we also consider the case of deterministically disabling CSMA. We run each of the three configuration for 10 times, and compare them in terms of (average) throughput and PRR.

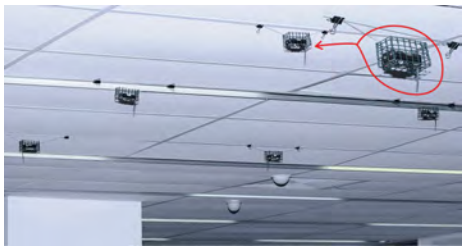

Fig. 12. A MicaZ-based WSN testbed on the ceiling of our research center.

As demonstrated by Fig. 13 (a), ART delivers a much higher throughput than the CSMA-based *BuzzBuzz*. Especially when $\omega = 500$, the throughput is nearly 5 times of the one delivered by *BuzzBuzz*, very close to the maximum achievable throughput at this frequency (indicated by the red dashed lines). Fig. 13 (b) shows that all transmission

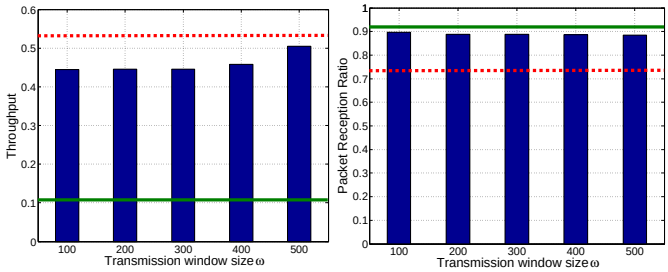

(a) Throughput (b) PRR

Fig. 13. The throughput (a) and PRR (b) of a transmission link at 2455 MHz with different transmission window sizes. The green solid lines indicate the performance of CSMA-based *BuzzBuzz*, while the red ones represent the performance of deterministically disabling CSMA.

window sizes yield a PRR of almost 0.9.This satisfies our target qualification range and goes very close to *BuzzBuzz* (which has a PRR of 0.92), even though *BuzzBuzz* employs sophisticated techniques (i.e., multiple headers and Reed-Solomon coding) to pursue perfect PRR based on CSMA. In a nutshell, ART obtains a considerable improvement in terms of throughput with a negligible compromise in PRR.

Fig. 13 also shows that, with an increasing transmission window size, throughput is slightly increased, while PRR experiences a minor degradation. As increasing ω results in less transmission windows within the same period, the reduced inter-window transmissions (all with CSMA enabled) then helps to improve the throughput to some extent. Moreover, less updates of the CSMA probability lead to poorer adaptability to the dynamic WiFi interference, which accounts for the minor degradation of PRR.

4.3 A Five-Link Scenario

We test a scenario with five links $\{l_i = (s_i, r_i), i = 1, ..., 5\}$ as shown in Fig. 14. We deliberately deploy the five links in a relatively small area (about 20 m^2), in order to emulate a small part of a densely deployed WSN (which we do not have in full scale at our disposal). The frequency allocations based on different schemes are shown in Table 1. It is obvious that, in the three schemes other than ART-FAVOR, the number of available channels is smaller than the number of links; therefore, some links do have to share the same channel, while our FAVOR utilizes the spectrum more efficiently and thus can allocate different channels (center frequencies) to the five links. Furthermore, FAVOR allocates as different frequencies as possible to the links that are closed to each other in distance. Finally, due to the density function (Fig. 11) used by FAVOR, three of five links (i.e., l_2, l_3 and l_5) choose frequencies that are less interfered by WiFi (thus with higher values in $\Phi(f)$). We believe that the superiority of ART-FAVOR can become more evident if more links are involved.

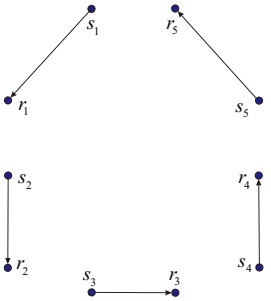

Fig. 14. A five-link deployment.

TABLE 1
Frequency allocations for the five-link scenario based on different schemes.

Links	l_1	l_2	l_3	l_4	l_5
Single-channel	2455	2455	2455	2455	2455
Two-channel	2450	2455	2450	2455	2450
Three-channel	2450	2453	2456	2450	2453
ART-FAVOR	2456	2453	2450	2458	2451

We test all seven schemes (with different combinations of channel allocations and CSMA control strategies), and report the results in Fig. 15 (a) and (b), respectively. For each reported data point, we perform 10 experiments, and plot their statistical quantities such as means, median, and interquartile ranges. As shown in Fig. 15(a), ART operates l_2, l_3 and l_5 much better than the other schemes: the three links nearly achieve their full data rates. As for l_1 and l_4 that suffer from WiFi interference, ART also performs fairly well as we expect: both links have a throughput much higher than the ones achieved by the schemes with CSMA deterministically enabled. The superiority of ART in throughput can also be confirmed by its very small interquartile ranges that actually indicate a very stable performance (which normally can be obtained only by enabling CSMA). In terms of PRR, Fig. 15(b) shows ART offers qualified PRRs for all five links, comparable to other schemes with CSMA enabled (whereas those schemes achieve an excellent PRR performance at the cost of poor throughput). In summary, as ART combines FAVOR and P-CSMA to exploit both frequency and temporal white space left out by WiFi interference, it offers opportunity to greatly improve throughput while maintaining the transmission quality.

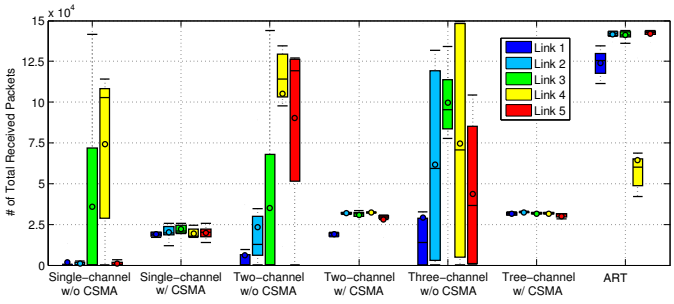

(a) Statistics on TRPs for all links

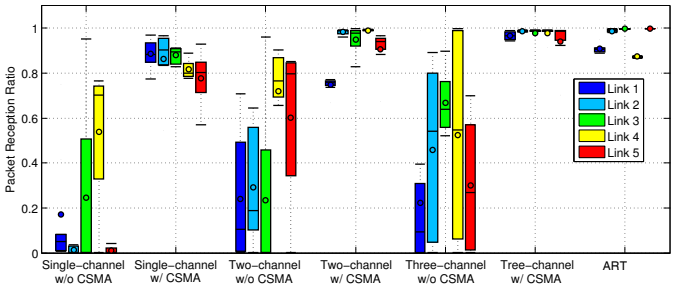

(b) Statistics on PRRs for all links

Fig. 15. Experiment results for the five-link scenario.

4.4 Tree-Based Data Collection

In order to demonstrate the benefit of applying ART to a practical application, we test the performance of different schemes in several arbitrary data collection trees deployed in a research center (an area of roughly $800m^2$). The deployment scenario is complicatedly structured with walls, such that the communications among sensor nodes may be obstructed. As their results are similar, we hereby report only one typical example. The node locations, as well as the network topology, are shown in Fig. 16 (a). It is shown that the sensor nodes are individually deployed in three different rooms, which results in multiple contention domains (i.e., the subsets of nodes in the same room). As proposed in Sec. 3.2.1, we need to perform our algorithm in each of them. The frequencies are allocated to non-leaf nodes to receive data packets from their child nodes. As we have discussed in Sec. 3.4, we apply the labeling method [20] to schedule links sharing the same incident node.

As shown in Sec. 4.3, the CSMA-disabled schedule needs to prevent the links with the same frequency from transmitting simultaneously. The more frequencies we can allocate, the less time slots we need in one round (during which every node gets a chance to transmit). As a result, two- and three-channel schemes need at least 29 and 24 slots respectively, while ART-FAVOR needs only 22 slots. For single channel, we enable CSMA and apply CTP [8] to perform data collection. We also enable CSMA for two- and three-channel schemes, but with a different schedule: it needs only to guarantee a sender and a receiver staying at the same frequency when the link between them is active. Consequently, only 9 time slots are needed. Given limited space, we show only the frequency allocation and schedule for FAVOR in Figure 16 (b).

We report the results in terms of TRP and PRR in Fig. 17 (a) and (b) respectively. Different from Sec. 4.3, here both

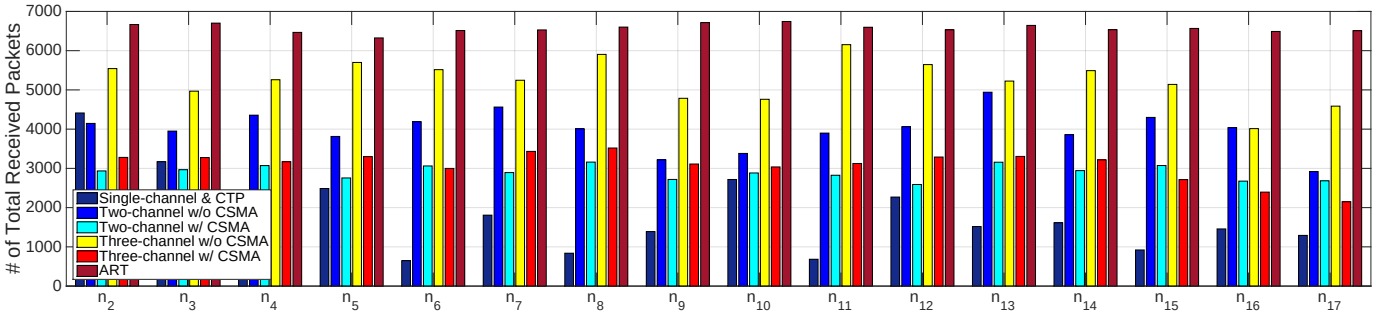

(a) Total Packet Reception (TPR) at the sink.

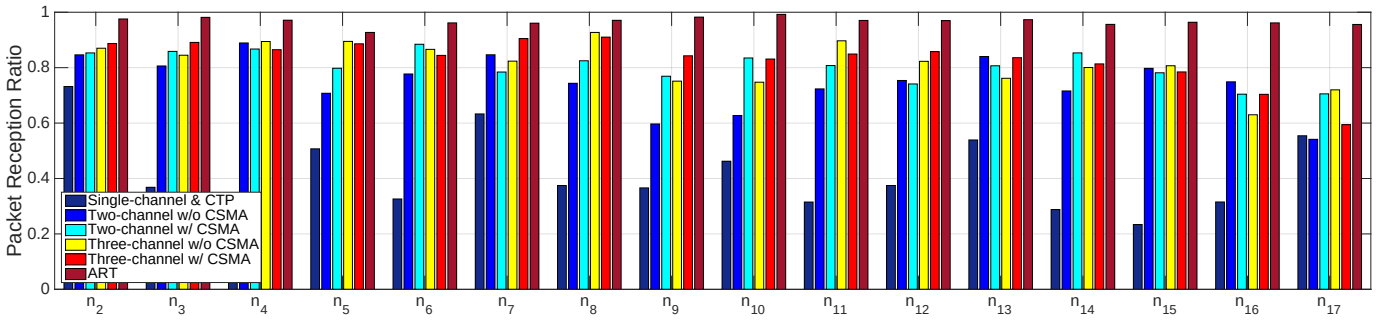

(b) Packet Reception Ratio (PRR) at the sink.

Fig. 17. Performance of six different schemes in a data collection tree.

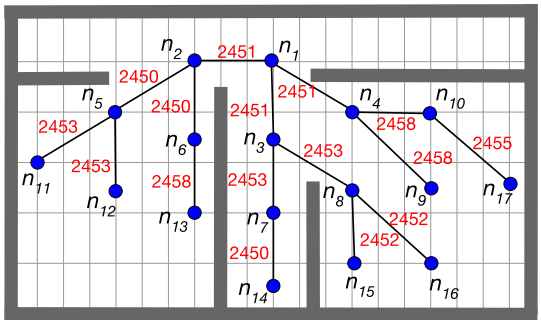

(a) Data collection tree (frequency allocation marked on links).

1	2	3	4	5	6	7	8	9	10	11	12
$n_2 \rightarrow n_1$						$n_3 \rightarrow n_1$					
$n_{14} \rightarrow n_7$	$n_{15} \rightarrow n_8$	$n_{16} \rightarrow n_8$				$n_{17} \rightarrow n_{10}$					

13	14	15	16	17	18	19	20	21	22
$n_4 \rightarrow n_1$				$n_5 \rightarrow n_2$				$n_6 \rightarrow n_2$	
$n_{11} \rightarrow n_5$	$n_{12} \rightarrow n_5$	$n_{13} \rightarrow n_6$		$n_7 \rightarrow n_3$		$n_9 \rightarrow n_4$	$n_8 \rightarrow n_3$		
							$n_{10} \rightarrow n_4$		

(b) Minimum delay schedule on the tree.

Fig. 16. Data collection tree with FAVOR frequency allocation (a) and min-delay transmission schedule (b). The node locations in (a) are plotted roughly proportional to their actual locations, and the thick gray lines denote walls.

TRP and PRR are defined for each node and measured at the sink. As expected, the results in Fig. 17 (a) show that ART surpasses all other schemes, and it achieves a throughput three times higher than the one reached by the single-channel CTP. Apparently, ART beats two- and three-channel schemes due to the less time slots in a round offered by FAVOR. Meanwhile, it prevails against all CSMA-enabled schemes thanks to P-CSMA's adaptive CSMA adjustment. In terms of PRR, it is shown in Fig. 17 (b) that ART achieves a nearly perfect transmission quality, even better than those CSMA-enabled schemes. The reason is that, when multiple co-channel transmissions proceed simultaneously, even CSMA cannot perfectly resolve such conflicts with its procrastination; whereas ART employs FAVOR to fully exploit the spatial-frequency space, providing sufficient room to counteract the interference between concurrent transmissions. This also reduces the dependence on CSMA and thus endows P-CSMA with more opportunities in "relaxing" CSMA to improve throughput. Overall, we have proved that ART is able to guarantee both throughput and PRR even under WiFi interference.

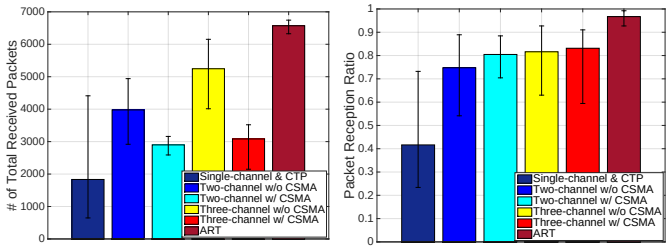

(a) Total Received Packets (TRP) (b) Packet Reception Ratio (PRR)

Fig. 18. Statistics of six different schemes in a data collection three.

To further reveal the performance of the whole network, we also report the average, maximum and minimum of TRP and PRR of the 16 sensor nodes under different schemes in Fig. 18. It is obviously shown that, our ART has the highest average TRP and PRR among all schemes. Furthermore, our ART has smallest performance variations, based on the fact that the maximum TRP (or PRR) of the sensor nodes is very

close to the minimum one. Hence, we conclude that, ART allows for a more fair sharing of the bandwidth; every node gets roughly the same transmission throughput and quality, which can be guaranteed by none of the other schemes.

4.5 Convergence of FAVOR

We briefly verify the convergence of FAVOR (in terms of communication rounds) in a 30-node WSN. Each node periodically communicates with its one-hop neighbors to calculate its receiving frequency, and also informs its one-hop neighbors of this temporary decision. We unify both frequency and distance in to a scale of $[0, 1]$ and employ the density function shown in Fig. 11 whose image is also normalized into $[0, 1]$. The A-CVT objective value is thus always small than 1. As illustrated by Fig. 19, the convergence under different step sizes (i.e., α) often takes 25-30 rounds. As such message exchanges can piggyback with other transmission activities and frequency (re)-allocation does not happen very often, the overhead of FAVOR (in terms of the entailed communication and computation costs) is affordable, especially considering the substantial improvement it is able to bring.

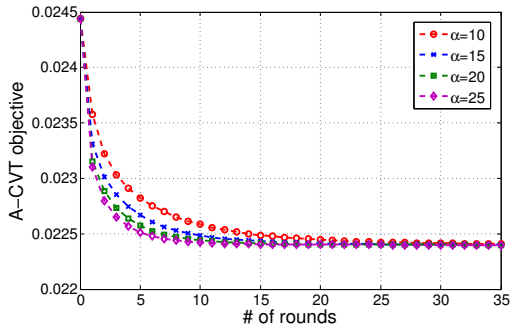

Fig. 19. The convergence of FAVOR.

5 Related Work and Discussions

Recent years have witnessed the development of multi-channel techniques to improve the performance of ZigBee-based WSNs (especially in terms of network throughput [12], [21]). Most proposals make use of graph-coloring heuristics to allocate channels in the whole network, but the work of [21] innovated in allocating channels to disjoint trees in WSNs. As the allocated channels are assumed to be orthogonal, the inter-channel interference is often neglected. However, the inter-channel interference does exist and its impact on link capacity and the performance of multi-channel protocols are systematically explored in [21], [23]. On the other hand, the existence of WiFi interference exerts more limitation to the availability of those orthogonal channels, as shown in Fig. 1.

To further improve the spectrum utilization in WSNs, partially overlapped channels have recently been introduced [25]. However, the proposal in [25] is still too conservative in using partially overlap channels, and it fails to discover the continuous frequency-distance relation in a spatial-frequency space. Moreover, all the aforementioned multi-channel schemes still keep ZigBee's default CSMA to handle possible transmission conflict, but as we have shown in Sec. 2.1, using CSMA to handle the internal ZigBee interference trades temporal resource for PRR thereby heavily restricting transmission throughput, while our FAVOR algorithm reaches not only much higher transmission throughput but also an almost perfect PRR.

Unfortunately, fully utilizing the whole ZigBee spectrum may cause certain ZigBee links to operate in the WiFi-interfering spectrum (which could be avoided should a single ZigBee channel, e.g., channel 25 or 26, be used). This co-existing issue has provoked certain interests from the research community [10], [14], [18], [24]. The proposals try to handle WiFi interference from every different perspectives. [18] proposes to model the correlation of interference and link burstiness in an off-line manner. Being aware of this model, ZigBee nodes could estimate the duration of the interference and then schedule the on-going transmissions by intelligently setting the backoff delay. [10] observes that WiFi frames are usually highly clustered in time, which results in a large number of temporal white spaces. Exploiting the self-similarity of the arrival process of WiFi frame clusters to predict the lengths of the white spaces, the data frames from MAC layer can be divided into sub-frames of smallers size and then scheduled into those "niches". [14] presents Multiple-Header (MH) mechanism, based on the observation that the header has the highest probability to be corrupted in a packet. [14] also employs Reed-Solomon coding to help correcting bit errors in packets. These two FEC mechanisms are deployed to complement the existing ARQ component in ZigBee, adding on a considerable overhead. [24] jointly takes both interference intensity and density into account, and designs a new method for assessing WiFi interference in a certain channel through periodically sampling the RSSI reading. By assessing the WiFi interference in all available channels, the nodes pick up the least interfered channel for receiving packets.

All the above existing proposals aim at enhancing the transmission quality but overlook the need for high throughput. To pursuing a perfect PRR, CSMA is unanimously adopted as a foundation for handling both internal ZigBee and external WiFi interference. However, as we have showed, CSMA heavily trades throughput for higher PRR. Moreover, the adaptations proposed by existing proposals introduce a large amount of extra overhead, which further jeopardize network throughput.

6 Conclusion and Future Work

In this paper, we have presented ART as a novel framework for guaranteeing both network throughput and transmission quality (in terms of PRR) for ZigBee-based WSNs in face of WiFi interference. ART consists of two components: FAVOR and P-CSMA.

Shifting the paradigm from discrete channel allocation to continuous frequency allocation, FAVOR innovates in transforming the optimal frequency allocation problem to a spatial tessellation problem in a frequency-spatial space, so as to adapt the CVT method to search for a near-optimal frequency allocation. By continuously exploiting the spectrum, FAVOR is able to utilize the frequency white space of WiFi interference; thereby significantly improving network

throughput. To cope with the survived WiFi interference, CSMA is employed in a smarter way. Aiming to maintain PRR at a satisfactory level, our P-CSMA is built upon CSMA but with a randomized control and a real-time interference assessment scheme, such that the use of CSMA can be adaptively tuned to exploit the temporal white space of WiFi interference. In order to demonstrate the efficacy of ART, we have performed intensive experiments on our MicaZ-based testbed under various network settings. All the experiment results have strongly proved the feasibility and superiority of ART in tackling the ZigBee-WiFi co-existence issue.

REFERENCES

[1] Chipcon's CC2420 2.4G IEEE 802.15.4/ZigBee-ready RF Transceiver.

[2] IEEE 802.15.4 1999. IEEE 802.15.4, Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Network (LR-WPANS). IEEE Std. 802.15.4, 2003.

[3] D. Bertsekas. *Nonlinear Programming*. Athena Scientific, Belmont, Massachusetts, 2 edition, 1999.

[4] C. Burnikel, K. Mehlhorn, and S. Schirra. How to Compute the Voronoi Diagram of Line Segments: Theoretical and Experimental Results. In *Proc. of the 2nd ESA (LNCS 855)*, 1994.

[5] P. Corke, T. Wark, R. Jurdak, W. Hu, P. Valencia, and D. Moore. Environmental Wireless Sensor Networks. *Proceedings of the IEEE*, 98(11):1903–1917, 2010.

[6] S. Dawson-Haggerty, S. Lanzisera, J. Taneja, R. Brown, and D. Culler. @scale: Insights from a Large, Long-Lived Appliance Energy WSN. In *Proc. of the 11th ACM/IEEE IPSN*, 2012.

[7] Q. Du, V. Faber, and M. Gunzburger. Centroidal Voronoi Tessellations: Applications and Algorithm. *SIAM Review*, 41(4):637–676, 1999.

[8] Omprakash Gnawali, Rodrigo Fonseca, Kyle Jamieson, David Moss, and Philip Levis. Collection Tree Protocol. In *Proc. of the 7th ACM SenSys*, 2009.

[9] J.-H. Hauer, V. Handziski, and A. Wolisz. Experimental Study of the Impact of WLAN Interference on IEEE 802.15.4 Body Area Networks. In *Proc. of the 6th EWSN*, 2009.

[10] J. Huang, G. Xing, G. Zhou, and R. Zhou. Beyond Co-existence: Exploiting WiFi White Space for ZigBee Performance Assurance. In *Proc. of the 18th IEEE ICNP*, 2010.

[11] J. Ko, J. Lim, Y. Chen, R. Musvaloiu-E, A. Terzis, G. Masson, T. Gao, W. Destler, L. Selavo, and R. Dutton. MEDiSN: Medical Emergency Detection in Sensor Networks. *ACM Trans. on Embedded Computing Systems*, 10(1):11:1–11:29, 2010.

[12] H.K. Le, D. Henriksson, and T. Abdelzaher. A Control Theory Approach to Throughput Optimization in Multi-Channel Collection Sensor Networks. In *Proc. of the 6th ACM/IEEE IPSN*, 2007.

[13] F. Li, J. Luo, G. Shi, and Y. He. FAVOR: Frequency Allocation for Versatile Occupancy of spectRum in Wireless Sensor Networks. In *Proc. of the 14th ACM Mobihoc*, 2013.

[14] C.-J.M. Liang, B. Priyantha, J. Liu, and A. Terzis. Surviving Wi-Fi Interference in Low Power ZigBee Networks. In *Proc. of the 8th ACM SenSys*, 2010.

[15] S. Lloyd. Least Square Quantization in PCM. *IEEE Trans. on Information Theory*, 28:129–137, 1982.

[16] A. Mishra, S. Banerjee, and W. Arbaugh. Weighted Coloring Based Channel Assignment for WLANs. *ACM Mobile Computing and Communication Review*, 9(3):19–31, 2005.

[17] Y. Shang and W. Ruml. Improved MDS-Based Localization. In *Proc. of the 23rd IEEE INFOCOM*, 2004.

[18] K. Srinivasan and M. Kazandjieva. The β-factor: Measuring Wireless Link Burstiness. In *Proc. of the 6th ACM SenSys*, 2008.

[19] L. Tang, Y. Sun, O. Gurewitz, and D.B. Johnson. EM-MAC: A Dynamic Multichannel Energy-Efficient MAC Protocol for Wireless Sensor Networks. In *Proc of the 12th ACM MobiHoc*, 2011.

[20] P.-J. Wan, Z. Wang, Z. Wan S.C. Huang, and H. Liu. Minimum-Latency Scheduling for Group Communication in Multi-channel Multihop Wireless Networks. In *Proc. of the 4th WASA (LNCS 5682)*, 2009.

[21] Y. Wu, J. Stankovic, T. He, and S. Lin. Realistic and Efficient Multi-Channel Communications in Wireless Sensor Networks. In *Proc. of the 27th IEEE INFOCOM*, 2008.

[22] T. Xiang, Z. Chi, F. Li, J. Luo, L. Tang, L. Zhao, and Y. Yang. Powering Indoor Sensing with Airflows: A Trinity of Energy Harvesting, Synchronous Duty-Cycling, and Sensing. In *Proc. of the 11th ACM SenSys*, 2013.

[23] G. Xing, M. Sha, J. Huang, G. Zhou, X. Wang, and S. Liu. Multi-Channel Interference Measurement and Modeling in Low-Power Wireless Networks. In *Proc. of the 30th IEEE RTSS*, 2009.

[24] R. Xu, G. Shi, J. Luo, Z. Zhao, and Y. Shu. MuZi: Multi-channel ZigBee Networks for Avoiding WiFi Interference. In *Proc. of the 4th IEEE CPSCom*, pages 323 – 329, 2011.

[25] X. Xu, J. Luo, and Q. Zhang. Design of Non-orthogonal Multi-channel Sensor Networks. In *Proc. of the 30th IEEE ICDCS*, 2010.

Feng Li received his BS and MS degrees in Computer Science from Shandong Normal University, China, in 2007, and Shandong University, China, in 2010, respectively. He got his PhD degree (also in Computer Science) from Nanyang Technological University, Singapore, in 2015. From 2014 to 2015, he worked as a research fellow in National University of Singapore, Singapore. He is currently an Assistant Professor at School of Computer Science and Technology, Shandong University, China. His research interests include computational geometry, applied optimization, distributed algorithms, wireless networking and mobile/pervasive computing.

Jun Luo received his BS and MS degrees in Electrical Engineering from Tsinghua University, China, and the PhD degree in Computer Science from EPFL (Swiss Federal Institute of Technology in Lausanne), Lausanne, Switzerland. From 2006 to 2008, he has worked as a post-doctoral research fellow in University of Waterloo, Canada. He is currently an Associate Professor at School of Computer Science and Engineering, Nanyang Technology University in Singapore. His research interests include wireless networking, mobile and pervasive computing, applied operations research, as well as network security.

Gaotao Shi received his Ph.D. degree from Harbin Institute of Technology, Harbin, Heilongjiang, China, in November 2006. He joined School of Computer Science and Technology, Tianjin University, Tianjin, China in December 2006. He was an Assistant Professor from December 2006 to June 2011. He is currently an associate Professor with Tianjin University since July 2011. His research interests include wireless sensor networks and embedded network systems.

Ying He received the BS and MS degrees in Electrical Engineering from Tsinghua University, China, and the PhD degree in Computer Science from the State University of New York (SUNY), Stony Brook, USA. He is currently an associate professor at the School of Computer Science and Engineering, Nanyang Technological University, Singapore. His research interests are in the broad areas of visual computing, with a focus on the problems that require geometric computation and analysis. More information can be found at http://www.ntu.edu.sg/home/yhe.