FAVOR: Frequency Allocation for Versatile Occupancy of spectRum in Wireless Sensor Networks

Feng Li* Jun Luo* Gaotao Shi† Ying He*

*School of Computer Engineering, Nanyang Technological University, Singapore

†School of Computer Science and Technology, Tianjin University, China

{fli3, junluo, yhe}@ntu.edu.sg, shgt@tju.edu.cn

ABSTRACT

While the increasing scales of the recent WSN deployments keep pushing a higher demand on the network throughput, the 16 orthogonal channels of the ZigBee radios are intensively explored to improve the parallelism of the transmissions. However, the interferences generated by other ISM band wireless devices (e.g., WiFi) have severely limited the usable channels for WSNs. Such a situation raises a need for a spectrum utilizing method more efficient than the conventional multi-channel access. To this end, we propose to shift the paradigm from discrete channel allocation to continuous frequency allocation in this paper. Motivated by our experiments showing the flexible and efficient use of spectrum through continuously tuning channel center frequencies with respect to link distances, we present FAVOR (Frequency Allocation for Versatile Occupancy of spectRum) to allocate proper center frequencies in a continuous spectrum (hence potentially overlapped channels, rather than discrete orthogonal channels) to nodes or links. To find an optimal frequency allocation, FAVOR creatively combines location and frequency into one space and thus transforms the frequency allocation problem into a spatial tessellation problem. This allows FAVOR to innovatively extend a spatial tessellation technique for the purpose of frequency allocation. We implement FAVOR in MicaZ platforms, and our extensive experiments with different network settings strongly demonstrate the superiority of FAVOR over existing approaches.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Network topology*

General Terms

Algorithms, Design, Performance

Keywords

Continuous Frequency Allocation, Wireless Sensor Networks, CC2420 Radio, Spatial Tessellation.

MobiHoc'13, July 29–August 1, 2013, Bangalore, India.

1. INTRODUCTION

Deeply exploited as an emerging technology, *Wireless Sensor Networks* (WSNs) have the potential to be widely deployed to support a variety of applications. In many recent applications, both the scale of WSN deployments and the demand in data rate for individual sensor nodes keep increasing [3,6]. However, such a development is severely hampered by the co-channel interference produced by the ever increasing wireless transmissions and their intensity. Moreover, co-channel interference may come from not only the ZigBee [2] devices involved in a WSN, but also other 2.4GHz ISM band occupants such as WiFi and Bluetooth [11, 12].

As ZigBee compatible radios (e.g., CC2420 of MicaZ Motes [1]) may operate on up to 16 channels, common wisdom suggests that one can make use of the multi-channel ability to prevent WSN links from interfering each other. This has led to quite a few research proposals, including prominently multi-channel scheduling for multi-hop transmissions (e.g., [24]) and multi-channel MACs (e.g., [21]). However, the number of channels available to a WSN is much lower than what ZigBee radios can offer, mainly due to the strong interference from other occupants in the 2.4GHz ISM band [11]. In our case (as shown in Figure 1), the experimental field is occupied by many WiFi testbeds, which effectively constrains the "clean" spectrum[1] to 2473–2483MHz where only two ZigBee channels can fit in [2].

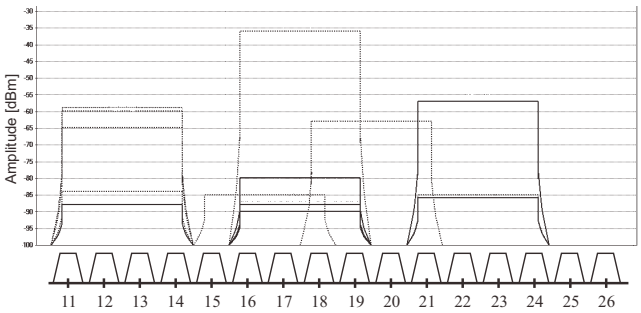

Figure 1: We overlap the WiFi signal strength observed through inSSIDer with the 16 channels of ZigBee. It is easy to spot that, apart from channels 25 and 26, other channels can be heavily interfered by multiple WiFi hotspots.

[1]Though avoiding WiFi interference temporally [12] or in frequency [28] is possible, directly incorporating these algorithms may severely complicate the system design, so we leave the ZigBee-Wifi coexistence issue as a future work.

By intensively experimenting (to be reported later) on our MicaZ Motes and their CC2420 radios, we have discovered that it would be more efficient to use the spectrum resource in a **continuous** manner. In other words, we should allocate channels to nodes or links in terms of their center frequencies (rather than a finite set of channel IDs). As it is well known that interference attenuates with distance and the distance between two interfering nodes/links is a **continuous** variable in Euclidean space, it is bounded to be more flexible and efficient (in term of spectrum utilization) to tune the frequencies of the nodes/links in a continuous manner, compared with the conventional graph coloring approach [18] where discrete channels are allocated based on the concept of (discrete) graph distance. Therefore, if one could allocate spectrum resource by jointly taking into account frequency and distance in a continuous manner, there is still a great potential to improve the performance of WSNs. For practical implementation, the ability of CC2420 radio in adjusting its channel center frequency at a granularity of 1MHz [1] does offer a good approximation of frequency fine-tuning.

Based on the above discovery, we first propose to shift the paradigm from *channel allocation* to (center) *frequency allocation* in utilizing frequency spectrum for WSNs. Although continuously allocating channel center frequencies to different links may lead to partially overlapped channels being operating simultaneously, we can combat the resulting interference by spacing them with a proper distance. In order to find the optimal frequency allocation for a set of spatially distributed nodes (or links), we further propose *Frequency Allocation for Versatile Occupancy of spectRum* (FAVOR) as a novel framework for frequency allocation in WSNs. FAVOR consists of two main components: i) a metric that unifies frequency and distance, which allows us to transform the frequency allocation problem into a spatial tessellation problem in a higher dimensional space, and ii) an algorithm that innovates on the *Centroidal Voronoi Tessellation* (CVT) method [8] to search for the local but nearly optimal frequency allocations.

Roughly speaking, FAVOR results in a frequency allocation such that nodes/links that are closer to each other in distance are further away from each other in frequency, while those far from each other in distance are allowed to be close in frequency. While 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 discrete sets to continuous spaces. The improved spectrum efficiency is bounded to favor WSN performance in throughput. Moreover, the FAVOR algorithm can be performed in a distributed manner using only local information. In order to verify the efficacy of FAVOR, we perform extensive experiments on a set of arbitrary links, as well as on two data collection trees. Our experiments demonstrate that FAVOR outperforms conventional graph coloring channel allocations and a recent overlapped channel allocation mechanism [29] in terms of throughput, given a stringent spectrum resource. In summary, our main contributions are:

- We propose to replace channel allocation with (center) frequency allocation, representing a paradigm shift in utilizing the scarce frequency spectrum.
- We define an optimization objective unifying frequency and distance; it enables us to formulate a frequency allocation problem into a spatial tessellation problem in a high dimensional space.
- We propose a new algorithm inspired by CVT to search for the (local) optimal frequency allocations; the algorithm entails an easy distributed implementation with a need for only local information.
- We perform extensive experiments, both to investigate the frequency-distance tradeoff and to verify the efficacy of our proposed FAVOR framework.

In the remaining of our paper, we first report the experiments that lead to our discovery of the distance driving frequency allocation method in Sec. 2. Then we present our FAVOR framework in Sec. 3. Further experiments in demonstrating the superiority of FAVOR are reported in Sec. 4. We also discuss the related work and possible extensions of FAVOR to general wireless networks in Sec. 5, before concluding our paper in Sec. 6.

2. MOTIVATION AND MATHEMATICAL BACKGROUND

In this section, we first motivate the design of FAVOR, by reporting our experiment results that exhibit an almost continuous tradeoff between (center) frequency and (Euclidean) distance for two competing wireless links. We also briefly discuss the mathematical background needed for the development of FAVOR optimal frequency allocation algorithm.

2.1 Frequency–Distance Tradeoff for 2 Links

Our simple experiment setting involves two parallel wireless links, l_1 and l_2, operated by four MicaZ nodes. The configuration is illustrated in Figure 2: we fix the location of transmission link l_1, and change the location of l_2 with the distance to l_1 ranging from 1.2 m to 4.8 m. We make the transmitters of the two links to send packets persistently as fast as possible, and measure the throughput of the two links, i.e., how many packets can be received correctly.

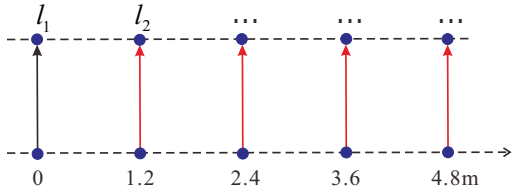

Figure 2: Experiment configuration. Each link has a length (the distance from its sender to its receiver) of 3.6m.

2.1.1 The Anomaly of MicaZ with CSMA

We first report an anomaly of MicaZ motes to better motivate our case.[2] It is well known that, though the *nominal data rate* of CC2420 is 250kbps, the maximum stable data rate one may squeeze from it is much lower (it is only about 50kbps for TelosB [22]). In our experience with MicaZ, we can only push one packet (45 bytes in total including 26 byes

[2]This anomaly also exists for other nodes using CC2420 radio, such as TelosB motes. However, we will not report results for other platforms due to the space limit.

pay load and 19 bytes header) through a ZigBee channel every 9ms, which achieves a data rate of 40kpbs (equivalent to a throughput of 0.16 with respect to the nominal data rate). In the following, we say a link achieving *full* data rate (or throughput) if it has a performance comparable to a single link. One would expect that operating two links simultaneously in the same channel would result in much lower (at least halved) data rate for each link. The results in Figure 3, nevertheless, show very different outcome, where we measure the throughput as the ratio between receiving rates at the receivers and the nominal data rate.

An amazing observation is that both links roughly achieve the full throughput of 0.16. The reason accounting for this

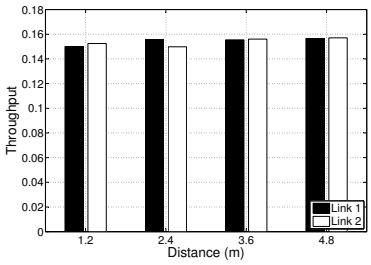

Figure 3: The throughput for two co-channel links with different distances.

anomaly is that the MCU (Atmel AVR Atmega 128L) used by MicaZ spends a certain amount of time moving packets within the protocol stacks and, most prominently, performing CSMA, the radio interface is left in an unsaturated mode. Consequently, CSMA may perfectly coordinate the two links such that they both achieve nearly full throughput. Actually, as we will show later, even adding up to 5 links will not drastically decrease the throughput of individual links. So one important message we have here is the follow:

> *Testing a channel allocation mechanism on CSMA-enabled platforms can lead to wrong conclusion, as even co-channel links may NOT conflict if we observe the outcome only in terms of throughput.*

2.1.2 Frequency vs. Distance without CSMA

In light of the results presented in Sec. 2.1.1, we test the throughput of two links under different (center) frequency and distance with CSMA disabled. While the four distances are shown in Figure 3, we also vary the frequency of l_1 from 2475MHz to 2480MHz but keep the frequency of l_2 at 2480MHz. When CSMA is disabled, a single link can carry a packet as fast as every 2ms, resulting in a data rate of 180kbps and a throughput (against the nominal data rate) of 0.72. Therefore, the full data rate in this case is 4.5 times higher than a CSMA-enabled link. The throughputs for all the frequency–distance combinations are shown in Figure 4. The figure clearly shows a throughput tradeoff surface we may achieve by extending from a 1D discrete frequency space (for conventional channel allocations) to a higher dimensional *spatial-frequency space*.

We also plot the two sets of results in 2D figures in Figure 5, with Figure 5(a) fixing the distance at 1.2m and Figure 5(b) fixing the frequency of l_1 at 2479MHz. It is clear from Figure 5(a) that even a frequency interval of only 1MHz is rather usable at the minimum distance: while the higher

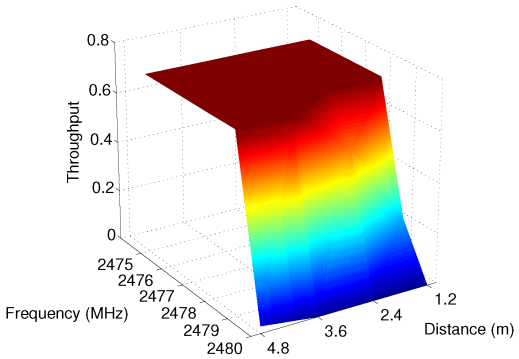

Figure 4: The lower throughput for two links with different frequency–distance combinations. CSMA is disabled for both links.

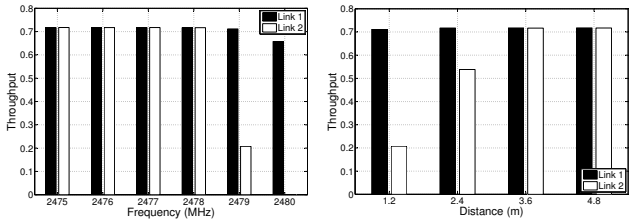

(a) Throughput vs. frequency (b) Throughput vs. distance

Figure 5: The throughput of both links as functions of frequency interval (a) and distance (b).

throughput is almost full (beyond 0.7), the lower throughput still offers a reasonable data rate of about 50kbps (even higher than the full rate of a single link with CSMA enabled). Figure 5(b) further shows that, by increasing the distance between two links, both links may achieve a nearly full throughput. Remember that, if we follow the IEEE 802.15.4 standard, only two independent channels are available for the spectrum we can use (due to the heavy interference from WiFi hotspots, see Sec. 1), whereas we almost get six usable but overlapped channels by varying center frequencies and distances in a continuous manner.[3] Therefore, another important message we get is this:

> *Continuously tuning frequencies allocated to links with respect to the distances between them allows for a more flexible and efficient use of spectrum.*

Note that a byproduct is that CSMA can be disabled as soon as the center frequencies of all links are slightly misaligned, reducing overhead and hence further improving throughput.

Now the question is, *given a set of wireless nodes or links (whose locations are already fixed), how can we find the optimal (center) frequency allocation for them?* Our FAVOR framework is exactly proposed to address this question. Before diving into the details of our proposal, we need to briefly discuss the mathematical background relevant to our algorithm designed for FAVOR.

[3]We are still confined by the current implementation of CC2420 radio, whose center frequency can be tuned only at a granularity of 1MHz. With more flexible radios developed in the future, we will have more channels available.

2.2 Mathematical Background on Spatial Tessellation and CVT

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.

For a fixed number of generators, varying their locations results in different tessellations of $\mathcal{A}$. One way to identify the "best" tessellation is to define a metric that measures the quality of a tessellation, and a typical metric is the "impact" of the generator u_i to a point v in its cell,[4] represented often by $\|v - u_i\|_{\ell^2}$ for $v \in V_i$. This ends up with an objective that represents the total impact of the generators to their individual cells:

$$\mathfrak{I}(\{A_i\}, \{u_i\}) = \sum_i \int_{A_i} \|v - u_i\|_{\ell^2}^2 \Phi(v) dv, \quad (1)$$

where $\Phi(v)$ indicates the density at location v. This objective is neither convex nor concave, so optimizing this objective may lead to many local minima (or maxima). According to the theory of *Centroidal Voronoi Tessellations* (CVT) [8], we have the following four basic conclusions:

- The Voronoi tessellation is optimal for a fixed set of points $\{u_i\}$.
- A (good) local optimal solution is when every u_i coincides with the gravity center (centroid) of its cell V_i.
- Lloyd's method [14] that moves each u_i to the centroid of its current cell in every iteration terminates at this local optimal solution with a linear convergence rate.
- Lloyd's method outputs $\{u_i\}$ that are uniformly distributed in $\mathcal{A}$ and, if $\mathcal{A} \subseteq \mathbb{R}^2$, the cells that are almost all regular hexagons.

The interesting observation of the outcome of CVT is that $\{u_i\}$ at termination are uniformly distributed in the space, which is intuitively related to our need of "spreading" nodes/links over the available frequency spectrum. In fact, our FAVOR framework is a non-trivial extension of CVT to a space involving both frequency and distance (or location).

3. FAVOR: A LOCATION-AWARE FREQUENCY ALLOCATION SCHEME

In this section, we first discuss our system model, then we present our FAVOR framework in terms of the optimization objective and the algorithm to find a local optimal solution. Finally, we discuss how the algorithm can be implemented in practical scenarios.

[4]There are many interpretations to this metric [8], we choose the one that is relevant to our design later.

3.1 System Model

We assume a WSN consisting of a set of sensor 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 readily extensible to 3D volume deployments in theory, we confine our scenarios to 2D deployments due to the limitations imposed by our experimental conditions. Let $\{u_i\}_{i=1,...,N}$ be the locations of the sensor nodes, where $u_i \in \mathbb{R}^2$. Given a frequency band $\mathcal{B} = [f_{min}, f_{max}]$ and the channel width f_w, we assign each sensor node n_i a channel with center frequency $f_i \in \mathcal{B}'$, where $\mathcal{B}' = [f_{min} + \frac{f_w}{2}, f_{max} - \frac{f_w}{2}]$.[5] Combining the node's location and frequency, a node n_i now has a new "coordinate" $(u_i, f_i) \in \mathbb{R}^2 \times \mathcal{B}'$. We denote by $\mathcal{N}(n_i)$ the one-hop neighbors of node n_i: the nodes with whom n_i can communicate directly given a common channel and a fixed transmit power.

3.2 FAVOR Objective: Balancing Distance and Frequency

According to our observation in Sec. 2.1.2, a good frequency allocation scheme should assign very different frequencies to nodes that are close to each other but arbitrary frequencies to nodes that are far from each other. We illustrate such a possible location-dependent frequency allocation scheme in Figure 6. However, as there are many possible allocations given a certain WSN deployment, we need to find the best possible allocations, for which we need a metric to evaluate different allocations.

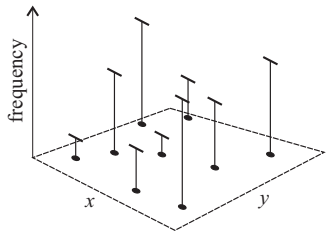

Figure 6: Allocating center frequencies based on the nodes' locations. While the black points indicate the sensor nodes, the "poles" on the points (with their different heights) represent different center frequencies allocated to the nodes.

To this end, we define the *impact metric* as

$$\|v - u_j\|_{\ell^2} + \beta \|f - f_i\|_{\ell^2},$$

where $v \in \mathbb{R}^2$ and $f \in \mathcal{B}'$, β is a weight to tune the tradeoff between frequency and distance, and we aim at a "tessellation" of the subset $\mathcal{A} = \mathbb{R}^2 \times \mathcal{B}'$ in $\mathbb{R}^3$ such that the following FAVOR objective is minimized:

$$\mathfrak{F}(\{A_i\}, \{f_i\}) = \sum_i \int_{A_i} \left(\|v - u_i\|_{\ell^2}^2 + \beta \|f - f_i\|_{\ell^2}^2\right) dz, \quad (2)$$

where $z = (u, f) \in \mathcal{A}$. It is obvious that (2) differs from (1) mainly in that part of the coordinates are fixed: we do not get the freedom to move nodes around, only the frequencies allocated to them are variables. Intuitively, $\mathfrak{F}(\{A_i\}, \{f_i\})$ represents the total "impact" from nodes to their cells. Given

[5]Channels can also be assigned to links rather than nodes; we will discuss this later in Sec. 3.4.

fixed locations, minimizing the objective has the effect of minimizing interference to some extent: according to what we discussed in Sec. 2.2, an optimal solution tends to spread out the generators (nodes' locations in this case). Since the locations in $\mathbb{R}^2$ are fixed, frequencies for two close-by nodes will be pushed further from each other. Another (relatively minor) difference is that we let $\Phi(v) = 1$ in (2), as both Euclidean and frequency spaces are assumed to be homogeneous for now. We could make use of $\Phi(\cdot)$ to characterize the non-homogeneity in space for our future development.

3.3 FAVOR: A CVT-based Approach

Given the similarity between (1) and (2), it is natural that one would propose to apply CVT to find a local optimal solution. Without loss of generality, we let $\beta = 1$ to simplify our derivation. The problem 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 Euclidean space), we need an algorithm more efficient than the Lloyd's method [14]; otherwise nodes with limited computation resource cannot afford it,[6] and ii) part of the coordinates for each $u_i \in \mathcal{A}$ are fixed, whereas CVT requires all the coordinates to be variables. To tackle these issues, we first propose *Approximate CVT* (A-CVT) to transform the problem into a more tractable and implementable form, then we 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 (2) as the following.

$$\mathfrak{F}(\{A_i\}, \{f_i\}) = \int_A \min_i \left(\|v - u_i\|_{\ell^2}^2 + \|f - f_i\|_{\ell^2}^2 \right) dz \quad (3)$$

The equivalence between (2) and (3) 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 shown in Sec. 2.2). Now we get a global integration over $\mathcal{A}$, eliminating the need for re-computing the Voronoi tessellations in every iteration, but we have to face the non-differentiable function $\min(\cdot)$. To make the problem tractable, we apply an approximation to the $\min(\cdot)$ function to "smooth" it, leading to the following A-CVT objective.

$$\mathfrak{F}(\{A_i\}, \{f_i\}) = \int_A \left[\sum_i \left(\|v - u_i\|_{\ell^2}^2 + \|f - f_i\|_{\ell^2}^2 \right)^{\lambda} \right]^{\frac{1}{\lambda}} dz \quad (4)$$

Due to page limit, we omit the proof of (4) converging to (3) when $\lambda \to -\infty$. In practice, we take $\lambda \in [-40, -20]$.

As minimizing the A-CVT objective (4) 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**. Roughly, the algorithm proceeds in rounds and takes the following three steps in each round:

1. Compute gradient for each generator z_i as (line 2)

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

   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 z 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 $\lambda \to \infty$. 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 WSN, we can use the communication neighborhood to approximate the tessellation neighborhood.

2. As u_i is fixed and we can change only f_i, we take $g_f(z_i)$ as the projection of $g(z_i)$ on the frequency axis (line 3).

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 8); otherwise, the frequency of each n_i is updated by $f_i \leftarrow f_i^+$, the outcome is exchanged among neighboring nodes (line 6), and further compuation will be conducted during the next round.

Algorithm 1: FAVOR

Input: For each $n_i \in \mathcal{N}$, location u_i, initial frequency $f_i^0 \in \mathcal{B}'$, stopping tolerance ε_1 and ε_2
Output: f_i^* for each n_i

1 For every node $n_i \in \mathcal{N}$ in each round (every τ ms):
2 Compute the gradient $g(z_i)$ of (4)
3 Project $g(z_i)$ on f_i to get $g_f(z_i)$
4 $f_i^+ = f_i - \alpha \cdot g_f(z_i)$ /*α is the step size*/
5 **if** $|g_f(z_i)| > \varepsilon_1 \vee |f_i^+ - f_i| > \varepsilon_2$ **then**
6 $\quad f_i \leftarrow f_i^+$; BROADCAST($f_i$) to nodes in $\mathcal{N}(n_i)$
7 **else**
8 $\quad f_i^* \leftarrow f_i^+$

We omit the convergence analysis as it follows directly from the basic theory of gradient-descent methods [4]. The convergence can be even faster if a centralized Quasi-Newton method is used to solve (4). We illustrate the results of our FAVOR algorithm in Figure 7. It is shown that, while the frequency allocation is initially ascending from left to right (with small frequency separations between neighboring nodes), the outcome of FAVOR exhibits much better separation in frequency for nodes close to each other. To facilitate visual illustration, we use only a 1D deployment (nodes on a line) as an example, which leads to easily discernable Voronoi tessellations in $\mathbb{R}^2$. However, our FAVOR algorithm works for any dimension higher than 2, while our experiments in Sec. 4 will be done for 2D deployments.

3.4 FAVOR for Disjoint Link Set

As FAVOR relies on a set of (point) locations $\{u_i\}$ to perform allocation, an obvious difficulty it may face is what if no obvious points exist in a networking scenario. In particular, a network may consist of several point-to-point wireless

[6] Lloyd's method requires to recompute the Voronoi tessellation in every iteration (due to the modified locations of the generators and the need to find the centroids of the cells), entailing high computational cost in $\mathbb{R}^n$ for $n > 2$. In fact, the complicated data structures required to model a $\mathbb{R}^3$ Voronoi tessellation cannot be easily implemented in sensor nodes, and no algorithm for CVT in $\mathbb{R}^n, n > 3$ has been implemented even for common CPUs.

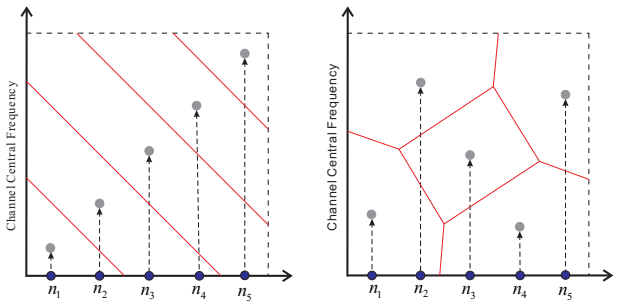

(a) Initial frequency allocation (b) Outcome of FAVOR

Figure 7: Frequency allocation based on Voronoi tessellations. With a regular (ascending) initial allocation (a), our FAVOR algorithm outputs a much better allocation (b). In both cases, the corresponding Voronoi tessellations are also shown using red line segments as cell boundaries.

links [29]. We propose two possible solutions for centralized and distributed computing separately. If a centralized computing is feasible, we may apply the extended Voronoi diagram where generators are not points but line segments (representing the links) [5]. Whereas this method may result in rather accurate frequency allocation, computing Voronoi tessellation with line segments as generators 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. If a point is chosen as the source or the destination, the computation is performed by that node. If the middle point is chosen, the computation can be done by either of the two nodes. Our experiments reported in Sec. 4.2 show that replacing line segments by points still leads to very good performance in throughput.

3.5 FAVOR For Tree-based Data Collection

One typical communication pattern in WSNs is the convergecast. More specifically, nodes of a WSN are organized into a data collection tree with a root at sink n_s, and every other node n_i sends data directly or indirectly (through multi-hop routing) to n_s [9]. In this case, frequency allocation itself is not sufficient to tackle the conflicts in media access: it cannot avoid conflicts for either multiple links ending at the same node or one node having both incoming and outgoing links, as such conflicts are the consequence of equipping a node with only one radio. Therefore, we need to perform both the frequency allocation and a TDMA-like time schedule. While a joint frequency allocation and scheduling problem is beyond the scope of our paper, we simply adapt the minimum latency scheduling mechanism [23].

We basically order the nodes into layers according to their hop-distances to the sink on the collection tree T, then use a similar labeling mechanism as proposed in [23]. 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 [23] adopts (orthogonal) channel allocation to enable parallel transmissions, our frequency allocation potentially allows more parallel transmissions. For the convenience of deriving bounds, channels are allocated using a first-fit distance-$(\rho + 1)$ graph coloring in [23], but we allocate different frequencies using FAVOR. As a node n_j needs to tune to the frequency f_i when transmitting to a receiver n_i, leaf nodes (being receivers to no one) do not need their own frequencies. We will show an example of our frequency allocation and scheduling in Sec. 4.3.

A similar frequency allocation and scheduling method also works for other types of communication patterns, such as broadcast (disseminating commands from the sink to the network) and aggregation (each relay node may send out less data then it receives). However, we focus only on data collection in our paper.

3.6 Time Synchronization

Time synchronization severs as a fundamental infrastructure for the TDMA scheduling used by FAVOR. Unfortunately, the existing synchronization protocols (e.g., FTSP [17]) rely on periodically flooding, resulting in a heavy traffic load that may significantly compromise the network throughput. In this paper, we employ a link-based approach that piggybacks control information with data traffic: for each transmission link, the transmitter synchronizes its transmitting schedule with that of the receiver based on the information piggyback with the acknowledgements sent by the receiver. To further suppress the control traffic, the transmitter may require acknowledgement only for the first transmission in each transmission schedule that consists of several transmissions as shown by Figure 11(b). In effect, non-root nodes report data to their respective parents while getting synchronized with the latter. Therefore, running this protocol within a data collection tree simply forces every node to stay synchronized with the root.

4. EVALUATION

We have implemented FAVOR in our MicaZ platforms and performed extensive experiments on them. The exciting experimental results are reported in this section. We first discuss the basic parameter settings for our experiments, then we describe our experiments on two different network scenarios. Finally, we briefly examine the convergence of FAVOR algorithm.

4.1 Experiment Settings

We apply MicaZ Motes and TinyOS 2.1 as the hardware and software platforms. As explained in Sec. 1, the available frequency spectrum we adopt is 2474–2481MHz, to avoid the "contaminated" spectrum by WiFi devices. According to the ZigBee standard [2], each channel has a 2MHz bandwidth and the center frequencies of two neighboring channels have to be separated by 5MHz. Therefore, only two orthogonal channels (with center frequencies 2475 and 2480) can fit into this spectrum (hence we term this scheme **two-channel**). However, we modify the codes for FAVOR to choose six possible center frequencies: 2475, 2476, 2477, 2478, 2479, and 2480MHz.[7] We will also test the proposal made in [29], where a center frequency interval of 3 or 2MHz has been suggested. This means that, given our available spectrum, three

[7]The CC2420 radio used by MicaZ Motes can adjust the frequency with only a granularity of 1MHz, which somewhat limits the potential of FAVOR. FAVOR can perform better if frequencies can be tuned continuously.

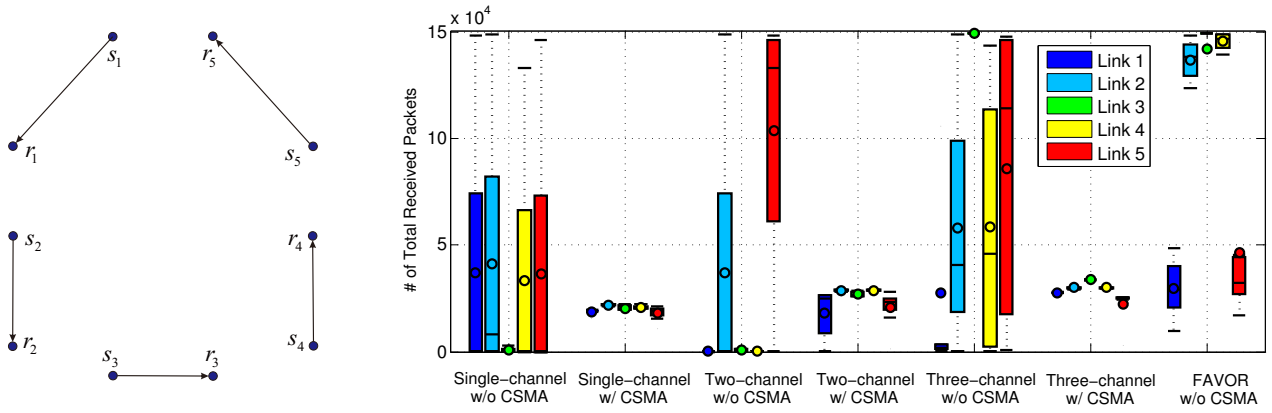

(a) Network deployment

(b) Statistics on TRPs for all links

Figure 8: A five-link scenario and the corresponding experimental results.

channels are available: 2475, 2477 (or 2478), and 2480MHz (hence we term this scheme **three-channel**). We apply a greedy graph coloring approach to allocate channels for both two- and three-channel schemes, while FAVOR uses its own algorithm for frequency allocation. In the following, we use only the center frequency to indicate a channel.

If we can avoid assigning the same channel to more than one link, we disable CSMA for all links. This allows packets to be pushed into a channel faster: one packet every 2ms according to Sec. 2.1.2. For cases where the same channel has to be assigned to different links, we test both with and without CSMA: we can only send one packet every 9ms in the former situation (see Sec. 2.1.2 again), whereas collisions may totally ruin some links in the latter situation, as already shown in Figure 5(a). We deploy our WSNs on the ceiling of our laboratory (see Figure 9), in order to emulate an indoor monitoring application scenario. We perform each experiment for 5 minutes, and use the *total received packets* (TRPs) at each destination node as the performance measure; it is actually an indicator of the **throughput**. For each reported data point, we perform 10 experiments, and plot their statistical quantities such as means and/or interquartile ranges.

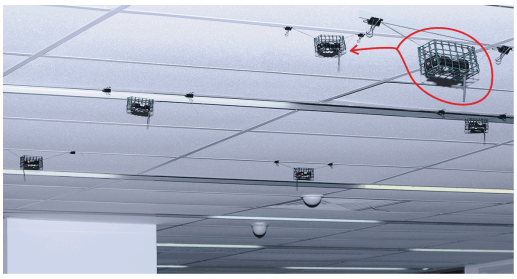

Figure 9: A MicaZ-based WSN testbed on the ceiling of our research center.

4.2 A Five-Link Scenario

We first test a scenario containing five links $\{l_i = (s_i, r_i), \forall i = 1, ..., 5\}$, as shown in Figure 8(a). We deliberately put the five links in a relatively small area (about $20m^2$), in order to mimic a small section of a densely deployed WSNs (which we do not have at our disposal). The optimal frequency allocations based on different schemes are shown in Table 1. Obviously, for schemes other than FAVOR, the number of available frequencies (channels) is smaller than the number of links, so some frequency has to be allocated to more than one link; the optimal allocation can simply try to space these co-channel links as far as possible.

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	2480	2480	2480	2480	2480
Two-channel	2480	2475	2480	2475	2480
Three-channel	2475	2477	2480	2475	2477
FAVOR	2476	2477	2480	2475	2478

We test all the four schemes: single-channel, two-channel, three-channel, and FAVOR; each with CSMA enabled and disabled (except FAVOR, as it does not need CSMA at all). The results in terms of TRPs per link are shown in Figure 8(b), where both interquartile ranges (boxes) and means (circles) are shown. It is obvious that FAVOR operates three out of five links (l_2, l_3, and l_4) much better than other schemes: mean TRPs can be up to three times higher than the one second to it. For another two links, FAVOR also does a relatively good job (better than any CSMA-enabled cases). Consequently, the total throughput of FAVOR is **five** times of the commonly used single channel with CSMA. Although three-channel without CSMA appears to achieve rather high mean TRPs in some links (l_3 and l_5), the huge interquartile ranges indicate a very unstable performance, rendering this scheme useless in practice.

The tradeoff between using CSMA or not (for other schemes) is very evident: CSMA delivers a rather stable performance by sacrificing throughput: the nature of a random access scheme. However, FAVOR achieves both stable and high throughput without the need for CSMA. In fact, FAVOR

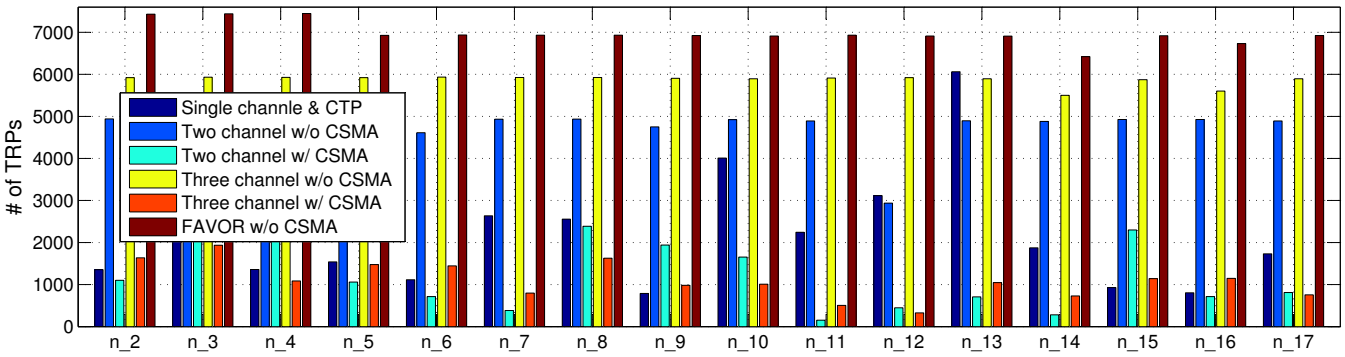

Figure 10: Performance of six different media access schemes in an unbalanced data collection tree.

could perform better if the radio allowed a finer granularity in tuning frequencies: we have to round the frequency allocation done by FAVOR to integer values of MHz, which leads to the relatively bad performance of l_1 and l_2.

4.3 Unbalanced Tree-based Data Collection

In order to demonstrate the benefit of applying FAVOR in a practical situation, we test the performance of different schemes in an arbitrary (unbalanced) data collection tree. The deployment area is roughly 600m^2, and the node locations, as well as the network topology, are shown in Figure 11(a). The frequencies (or channels) are allocated to non-leaf nodes, and we apply the labeling method in [23] for link schedules. Note that the disjoint-tree based scheduling [24] does not apply to this small network.

According to what is shown in Sec. 4.2, the schedule needs also to avoid links with the same frequency transmitting simultaneously. This leads to the obvious consequence that 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-channel scheme needs at least 29 slots, three-channel scheme needs at least 24 slots, while FAVOR needs only 22 slots. For single channel, we enable CSMA and apply CTP [9] 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 11.

As expected, the results in Figure 10 show that FAVOR surpasses all other schemes, and it achieves a throughput that is 3.36 times of that achieved by CTP. Apparently, FAVOR beats two- and three-channel schemes due to the less time slots in a round, and it prevails against all CSMA-enabled cases thanks to the elimination of the overhead brought by CSMA. Moreover, FAVOR allows for a more fair sharing of the bandwidth (every node gets roughly the same throughput), which cannot be guaranteed by any CSMA-based schemes.

4.4 Balanced Tree-based Data Collection

One may argue that FAVOR's continuous (hence potentially irregular) frequency allocation cannot offer significant network performance in a regular network topology, e.g., a balanced data collection tree. Therefore, we re-deploy

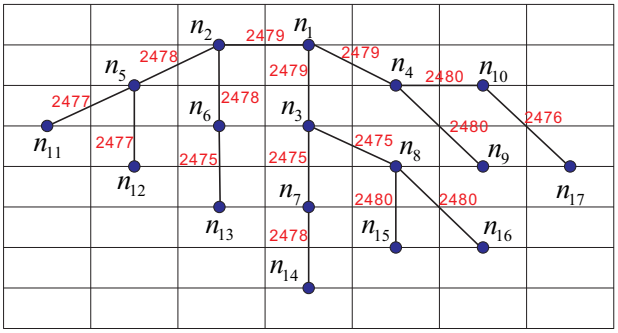

(a) Data collection tree (frequency allocation marked on links).

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.

Figure 11: Unbalanced 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 in our ceiling testbed.

our testbed to form a balanced data collection tree with nodes regularly spaced, as shown in Figure 12. All settings are maintained as those in Sec. 4.3, except that different FAVOR frequency allocation and transmission schedule are computed to suit this tree.

The results in term of TRPs are reported in Figure 13. Obviously, the observations made for Figure 10 still hold here, except that the advantage of FAVOR over others has been slightly reduced. This stems from the reduced number of links involved in this tree: as FAVOR surpasses others by offering higher frequency utilization, its advantage becomes more conspicuous if a higher utilization is actually needed.

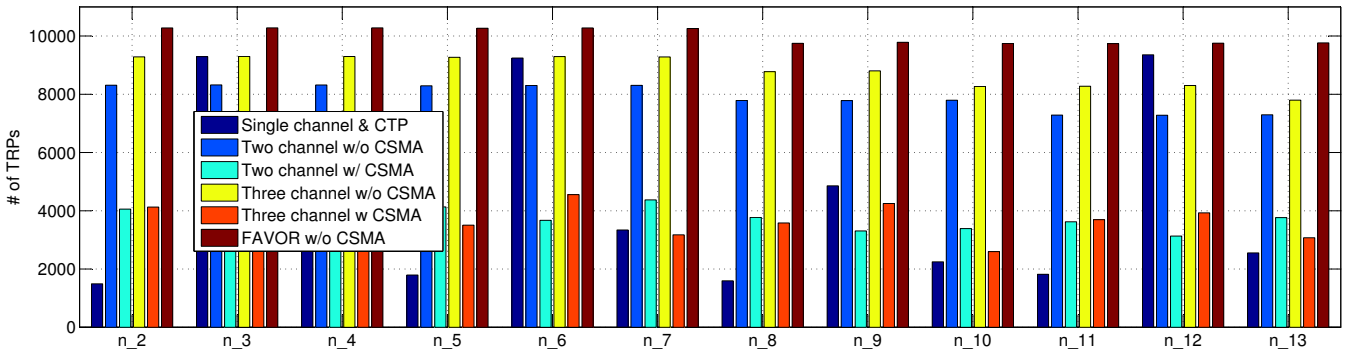

Figure 13: Performance of six different media access schemes in a balanced data collection tree.

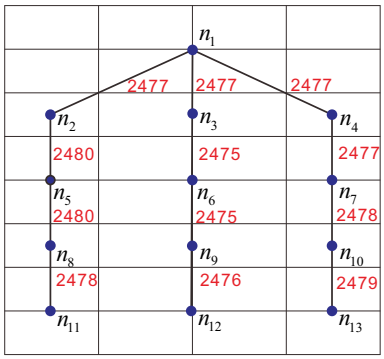

Figure 12: Balanced data collection tree with FAVOR frequency allocations marked on links. The node locations are plotted roughly proportional to their actual locations in our ceiling testbed.

4.5 Convergence of FAVOR

We briefly verify the complexity (in terms of communication rounds) of FAVOR in a distributed computing scenario: a 30-node WSN. As we unify both frequency and distance into a scale of $[0, 1]$, the A-CVT objective value is always smaller than 1. As shown in Figure 14, the convergence under different step size α's often takes 20–30 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, given the substantial throughput improvement it can bring.

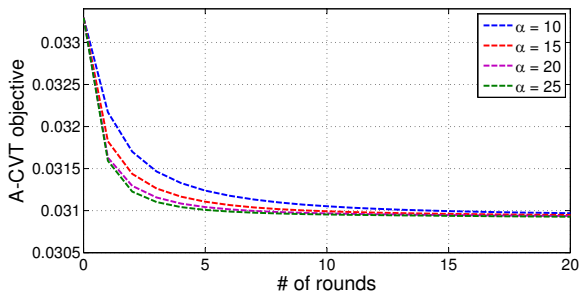

Figure 14: The convergence of FAVOR in 30-node WSN.

5. RELATED WORK AND DISCUSSIONS

While exploiting multi-channel access to improve the performance of general multi-hop wireless network has been intensively studied in the last decade (e.g., [20]), dedicated investigations for WSNs started rather late [13, 24]. Most proposals make use of graph-coloring heuristics to allocate channels in the whole network, but the work of [24] 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 [24, 27].

Given the limited number of non-overlapping (or orthogonal) channels, partially overlapped channels were later introduced to improve the spectrum utilization in WiFi networks [19]. However, it is only recently that the special inter-channel interference feature of ZigBee radio was identified [29]. In particular, whereas the center frequencies of two partially overlapped channels have to be sufficiently apart from each other for WiFi radios to properly operate [19], the ZigBee radios used by WSNs are far more robust due to their simple design: as we have also shown in Sec. 2.1.2, only a difference of 1MHz is enough for two channels to deliver a reasonable throughput. However, the work in [29] did not discover the advantage of continuous frequency allocation; it focuses only on adjusting CSMA. As demonstrated by our FAVOR, CSMA may not be needed anymore if frequency allocation is applied.

The multi-channel feature of ZigBee radios has also been applied to improve the quality of individual links [7] and to avoid the interference from WiFi devices [28]. However, the scalability of these proposals are still confined by the limited channels. We believe that the flexible spectrum utilization offered by FAVOR can also contribute to tackling these problems, as well as problems under a joint routing and link scheduling framework (e.g., [16]) or for duty-cycled WSNs (e.g., [10]).

Currently, FAVOR cannot be directly applied to WiFi networks due to the significant inter-channel interference between two partially overlapped channels [25]. This may stem from the particular filter design for a WiFi radio receiver, but we may be able to redesign the filter to reduce the inter-channel interference (hence to apply FAVOR for achieving a higher spectrum utilization), at a cost of slightly reduced data rate.

6. CONCLUSION

In this paper, we present FAVOR, a novel framework for efficient spectrum utilization in WSNs. We exploit a continuous frequency allocation to replace the conventional discrete multi-channel allocation. We then combine frequency and location into one space and thus transform the optimal frequency allocation problem into a spatial tessellation problem. Our FAVOR algorithm innovates on the Centroidal Voronoi Tessellation method to search for the nearly optimal frequency allocations. Finally, we perform extensive experiments to demonstrate the feasibility and superiority of FAVOR. For future work, we plan to apply FAVOR to broader scenarios including data aggregation WSNs (e.g., [26]) and WSNs with mobile elements (e.g., [15]).

Acknowledgments

We would like to thank the anonymous reviewers for their insightful comments and constructive suggestions. This work was supported in part by AcRF Tier 2 Grant ARC15/11 and Tianjin Research Program of Application Foundation and Advanced Technology (China) No. 12JCQNJC00200.

7. REFERENCES

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

[2] IEEE Standard 802.15.4, 2011.

[3] TOES: Terrestrial Ecology Observing Systems.

[4] D. Bertsekas. *Nonlinear Programming*. Athena Scientific, Belmont, MA, 2 edition, 1999.

[5] 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.

[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] M. Doddavenkatappa, M. Chan, and B. Leong. Improving Link Quality by Exploiting Channel Diversity in Wireless Sensor Networks. In *Proc. of the 33rd IEEE RTSS*, 2011.

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

[9] O. Gnawali, R. Fonseca, K. Jamieson, D. Moss, and P Levis. Collection Tree Protocol. In *Proc. of the 7th ACM SenSys*, 2009.

[10] K. Han, Y. Liu, and J. Luo. Duty-Cycle-Aware Minimum-Energy Multicasting in Wireless Sensor Networks. *IEEE/ACM Trans. on Netw.*, 21(3):910–923, 2013.

[11] 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.

[12] 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.

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

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

[15] J. Luo and J.-P. Hubaux. Joint Sink Mobility and Routing to Increase the Lifetime of Wireless Sensor Networks: The Case of Constrained Mobility. *IEEE/ACM Trans. on Netw.*, 18(3):871–884, 2010.

[16] J. Luo, C. Rosenberg, and A. Girard. Engineering Wireless Mesh Networks: Joint Scheduling, Routing, Power Control and Rate Adaptation. *IEEE/ACM Trans. on Netw.*, 18(5):1387–1400, 2010.

[17] M. Maróti, B. Kusy, G. Simon, and Á. Lédeczi. The flooding time synchronization protocol. In *Proceedings of the 2nd ACM SenSys*, pages 39–49, 2004.

[18] 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.

[19] A. Mishra, V. Shrivastava, S. Banerjee, and W. Arbaugh. Partially Overlapped Channels Not Considered Harmful. In *Proc. of the ACM/IFIP SIGMetrics/Performance*, 2006.

[20] J. So and N. Vaidya. Multi-Channel MAC for Ad Hoc Networks: Handling Multi-Channel Hidden Terminals Using a Single Transceiver. In *Proc. of the 5th ACM MobiHoc*, 2004.

[21] 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.

[22] P. Trenkamp, M. Becker, and C. Goerg. Wireless Sensor Network Platforms – Datasheets versus Measurements. In *Proc. of IEEE SenseApp*, 2011.

[23] 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 WASA*, 2009.

[24] 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.

[25] L. Xiang and J. Luo. Joint Channel Assignment and Link Scheduling for Wireless Mesh Networks: Revisiting the Partially Overlapped Channels. In *Proc. of the 21st IEEE PIMRC*, pages 2063–2068, 2010.

[26] L. Xiang, J. Luo, and C. Roserberg. Compressed Data Aggregation: Energy Efficient and High Fidelity Data Collection. *IEEE/ACM Trans. on Netw.*, 2013 (accepted to appear).

[27] 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 IEEE RTSS*, 2009.

[28] 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.

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