Bringing Multi-Antenna Gain to Energy-Constrained Wireless Devices

Sanjib Sur†, Teng Wei† and Xinyu Zhang
University of Wisconsin-Madison
{sur2, twei7}@wisc.edu and xyzhang@ece.wisc.edu
†Co-primary authors

ABSTRACT

Leveraging the redundancy and parallelism from multiple RF chains, MIMO technology can easily scale wireless link capacity. However, the high power consumption and circuit-area cost prevents MIMO from being adopted by energy-constrained wireless devices. In this paper, we propose Halma, that can boost link capacity using multiple antennas but a single RF chain, thereby, consuming the same power as SISO. While modulating its normal data symbols, a Halma transmitter hops between multiple passive antennas on a per-symbol basis. The antenna hopping pattern implicitly carriers extra data, which the receiver can decode by extracting the index of the active antenna using its channel pattern as a signature.

We design Halma by intercepting the antenna switching and channel estimation modules in modern wireless systems, including ZigBee and WiFi. Further, we design a model-driven antenna hopping protocol to balance a tradeoff between link quality and dissimilarity of channel signatures. Remarkably, by leveraging the inherent packet structure in ZigBee, Halma's link capacity can scale well with the number of antennas. Using the WARP software radio, we have implemented Halma along with a ZigBee- and WiFi-based PHY layer. Our experiments demonstrate that Halma can improve ZigBee's throughput and energy efficiency by multiple folds under realistic network settings. For WiFi, it consumes similar power as SISO, but boosts throughput across a wide range of link conditions and modulation levels.

Categories and Subject Descriptors

C.2.1 [**Computer Communication Networks**]: Network Architecture and Design—*Wireless Communications*; C.2.2 [**Computer Communication Networks**]: Network Protocols

Keywords

MIMO, Energy Efficiency, Mobile Devices

IPSN '15, April 14 - 16, 2015, Seattle, WA, USA.

http://dx.doi.org/10.1145/2737095.2737099.

1. INTRODUCTION

MIMO has been a key enabling technology for recent high-rate wireless standards. Compared with conventional SISO links, a MIMO transmitter can reduce bit-error-rate (BER) by redundantly coding the same data symbol through multiple antennas, thus achieving *diversity gain.* It also allows parallel transmission of different symbols through different antennas, thus achieving *multiplexing gain.* Both diversity and multiplexing mechanisms can scale throughput with the number of antennas without adding new spectrum.

However, a MIMO radio must accompany each antenna with a separate RF chain. Most components in the RF chain build on analog technologies that hardly benefit from Moore's law and remain fundamentally unchanged in the past two decades [1]. More critically, they account for the majority of the transceiver's power cost. Recent measurement studies revealed that MIMO power consumption increases linearly with the number of RF chains [2–4], which often nullifies the improvement in link capacity, resulting in even lower energy-per-bit than SISO. This is why most energy-constrained wireless devices, such as WiFi-equipped smartphones and ZigBee sensors, do not support MIMO.

Principle of Halma. In this paper, we propose a simple mechanism, called Halma,[1] that aims to bring multi-antenna benefits to energy-constrained wireless devices. The key idea lies in an antenna hopping scheme, inspired by the communication theoretic concept of space-shift keying (SSK) [5]. As illustrated in Figure 1, a Halma transmitter can run on a single RF chain, but it switches between multiple passive antennas, and uses the index of the antenna to convey extra bits of information on top of its original symbols. The receiver uses a single antenna. While decoding the original symbols, it can decipher the transmit antenna index inside each symbol. Different transmit antennas' symbols are distorted by the channel in different ways. The distortion can be modeled as a complex multiplier, which the receiver can use as a signature to track down the transmit antenna index.

Such a per-symbol antenna-hopping or SSK mechanism has been analyzed in information theory, and shown to improve link capacity *logarithmically* with the number of transmit antennas N_t [5]. But this assumes zero antenna-index decoding error, which in turn relies on diversity mechanisms from multiple RF-chains at receiver side, and consequently compromises energy efficiency [6]. In contrast, Halma focuses on achieving high energy efficiency. In particular, for

[1]Halma (from the Greek word "jump") is a board game where players strategically move pieces in sequence across a grid of squares.

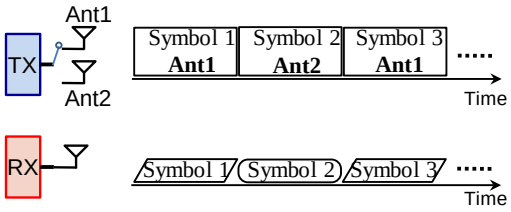

Figure 1: Example illustration of Halma. Transmitter: transmit each *data symbol* through a single-antenna, but hop between antennas on a per-symbol basis. Receiver: decode each data symbol, along with the transmitter's antenna index that implies extra bits of information.

a ZigBee link with single RF-chain transmitter and receiver, Halma can scale link capacity with N_t at a even faster rate than SSK, which translates into enormous energy saving.

Halma achieves this goal by uncovering the hidden potential of antenna hopping in real communications systems like ZigBee. It employs an *antenna index coding* (AIC) framework that enables a *fine-grained* antenna hopping. The key observation is that real wireless devices need to compound symbol-level modulation (*e.g.*, BPSK) with wide-band channel spreading (*e.g.*, DSSS). Consequently, Halma can embed multiple bits of antenna-index information in each original data symbol, by using *sub-symbol* level antenna hopping. Further, to obviate the need for multiple RF chains at the receiver side, Halma judiciously plants redundancy in the antenna hopping patterns, such that decoding error can be minimized without incurring too much overhead.

In addition, conventional SSK commonly adopts simplified channel fading (*e.g.*, Rayleigh/Rician) models between different transmit antennas and the receive antenna. Under such models, it is optimal to exploit all N_t antennas to maximize capacity. Our test experiments disprove such assumptions. Intuitively, employing more antennas allows more bits to be conveyed through antenna switching, yet it may increase the BER of original data symbols that is bottlenecked by the weakest channel. Halma employs an *adaptive antenna hopping* (AAH) protocol that efficiently selects the subset of antennas to optimize this tradeoff, based on a model-driven framework instrumented by channel profile measurement.

We show that the underpinning principles of AIC and AAH can work for not only single-carrier ZigBee modulation, but also multi-carrier WiFi OFDM. For WiFi, Halma performs antenna hopping in the frequency domain – across the OFDM subcarriers. This requires a multi-RF-chain transmitter, although the receiver still runs on a single RF-chain. Thus, Halma-WiFi is best applicable to the downlink of wireless LANs with energy-constrained clients.

Testbed validation. To validate the Halma design, we have implemented it on the WARP software radio platform. We first develop a DSSS and OFDM modulation/demodulation library following the 802.15.4 and 802.11n PHY-layer specifications. Then, Halma's AIC and AAH protocols are built on top of the library. Our experiments demonstrate that Halma can boost ZigBee's link rate by 4.7× with 4 TX antennas — a super-linear gain owing to its sub-symbol level antenna hopping. Meanwhile, the rate improvement translates into more than 50% of energy reduction under a variety of settings.

As for WiFi, Halma achieves around 30% throughput gain over SISO across a wide range of SNR conditions, and even outperforms 802.11n's diversity coding scheme. Due to its restriction of single RF-chain at receiver side, Halma cannot beat WiFi's MIMO spatial multiplexing mode in terms of throughput gain, but it consumes less energy-per-bit under practical traffic patterns.

Contributions. In summary, we make the following contributions through the Halma design:

(i) We innovate an antenna index coding (AIC) mechanism that overcomes the limitations of conventional SSK, and can achieve super-linear link capacity growth through fine-grained antenna hopping, which translates into substantial energy saving for low-power wireless devices.

(ii) We invalidate the greedy approach of employing all available antennas for SSK, identify a tradeoff between link quality and effectiveness of antenna index coding, and design an AAH protocol to make the optimal balance.

(iii) We implement Halma on top of a ZigBee/WiFi PHY layer, and verify its feasibility and effectiveness through extensive testbed experiments. All our implementation and experimental data have been made open-source [7].

The reminder of this paper is structured as follows. Section 2 investigates the energy cost of conventional MIMO and the feasibility of Halma's AIC. Section 3 details the design components of Halma. We then describe the implementation of Halma (Section 4), and conduct a comprehensive evaluation (Section 5). Section 6 discusses practical considerations (Section 6), followed by a survey of related work (Section 7). Finally, Section 8 concludes the paper.

2. MOTIVATION

In this section, we motivate Halma's single RF-chain design by examining the energy cost of existing multi-RF-chain MIMO WiFi/ZigBee. Then, we empirically explore the feasibility of Halma's antenna index modulation/decoding.

2.1 MIMO: the Energy Cost

Existing work measured the power consumption of 3×3 WiFi MIMO adapters with PCIe interfaces, including Atheros 9380 and Intel 5300 [2–4], which observed a linear growth of power consumption with the number of active antennas (RF chains). Here we further explore whether the phenomenon is present in a broader class of devices including: *(i)* a USB-powered WiFi MIMO adapter, Linksys AE3000, that supports 3×3 MIMO, and *(ii)* a ZigBee MIMO device, Atmel REB233SMAD, that can activate two receive antennas simultaneously [8].

For the former, we first use a USB extension cord to expose the interface between the AE3000 and its host PC, and then use the Monsoon power monitor [3] to intercept the power supplier circuit and perform the measurement. We also modified the AE3000 open-source driver so that the device can be fixed at a desired transmission mode and number of antennas. For contrast, we also monitor the Atheros 9380 and Intel 5300 cards using a PCIe extension cable. Figure 2 lists the power consumption under different settings, each value being the average within a 5-minute data collection.

The USB adapter's TX, RX and idle power consumption all grows linearly as the number of active antennas increases. In particular, with 3 antennas, the idle power is 1.3× that of 1-antenna case. Whereas for TX and RX mode, it is 2.2× and 2×, respectively. Analysis of real network traffic

Modes		Device power consumption (W)		
		Atheros 9380	Intel 5300	Linksys AE3000
Sleep		0.13	0.22	0.15
Rx Idle	1	0.68	1.27	0.84
	2	0.80	1.39	0.96
	3	0.94	1.53	1.10
Rx data	1	1.38	1.34	0.83
	2	1.42	1.48	1.31
	3	2.06	1.65	1.60
Tx data	1	1.44	1.44	0.87
	2	1.46	1.50	1.35
	3	2.09	1.99	1.92

Figure 2: Power consumption of state-of-the-art WiFi MIMO transceivers.

revealed that WiFi devices typically spend more than 80% of time in idle listening mode [9]. Although a 3×3 MIMO can reduce transmission time to $\frac{1}{3}$ compared with SISO, it does not reduce the idle listening time [4, 10]. Suppose TX (RX) time is 10%, then the energy cost per-bit compared with SISO is roughly: $1.3 \times 80\% + \frac{2.2}{3} \times 10\% + \frac{2}{3} \times 10\% = 1.2\times$. Thus, although the transmission cost is reduced to $\frac{1}{3}$, overall MIMO actually consumes more energy/bit than SISO. For PCIe devices, the power cost scaling differs slightly from USB adapter, whereas the increase of energy per-bit still holds especially for chatty traffic patterns [2–4].

Measurement of real MIMO WiFi networks also consistently showed their lower energy efficiency [2–4], even though the throughput grows linearly with number of active antennas. This explains why MIMO is commonly avoided by battery-powered WiFi devices, such as smartphones. Our Halma scheme is designed to overcome this barrier, which harvests throughput gain from multiple antennas, but without the formidable energy cost of conventional MIMO.

Note that, the 802.11n standard incorporates a Spatial Multiplexing Power Save (SMPS) mode that adaptively switches from multiple to single RF chain during idle listening mode, but requires tedious messaging overhead that reduce the effective throughput. Also, because of the current h/w limitations, this switching can not be performed on per-packet basis and thus leads to lower energy efficiency than SISO [11].

As for the multi-antenna ZigBee board, it can activate two receiving RF chains during packet header searching, but only the receive antenna with higher signal strength is used during actual packet reception, to provide diversity selection gain. We represent its 2×2 *receive* power consumption by monitoring the header searching mode using the WARP software-radio, and predict the 2×2 *transmit* power consumption based on its TX/RX power ratio in SISO. Figure 3 shows the power consumption of 2×2 MIMO TX/RX is roughly 2× that of SISO. Since ZigBee has a low duty-cycle and can operate in TDMA mode, less power is wasted in idle listening compared with WiFi. However, even assuming all its power is effectively used for transmission/receiving, MIMO's energy-per-bit would be comparable to SISO.

2.2 Feasibility of Halma

Halma uses transmit antenna index to implicitly carry extra bits of information, which the receiver can decipher by *using different TX antennas' channel distortion patterns as signature*. Intuitively, the effectiveness of such a scheme depends on: (i) the consistency of an antenna's channel signature across a packet and (ii) the dissimilarity of signatures across antennas. The former holds because the channel coherence time is much longer than typical WiFi/ZigBee packet duration for static/pedestrian scenarios [12]. To verify the latter, we leverage our implementation of 802.11 channel estimation module (described in Section 4) and compute the Euclidean distance between the channel signatures of 4 TX antennas separated 6 cm away from each other. The receiver is randomly placed within line-of-sight of the transmitter (which may adversely increase channel similarity) in an office environment. In each experiment, we keep collecting the TX antennas' channel signatures for 2 ms (roughly a WiFi/ZigBee packet duration). Figure 4 illustrates the confusion matrix between the collected antenna signatures.

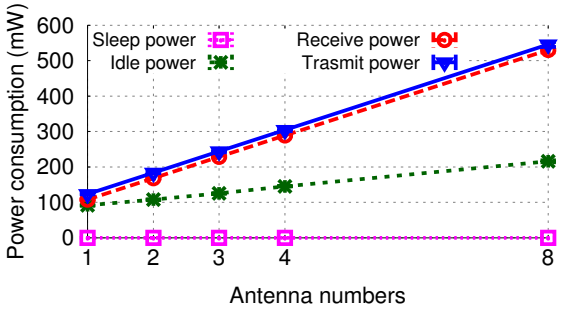

Figure 3: Power consumption of a multi-antenna ZigBee node. Cases with > 2 antennas are estimated using linear curve fitting.

Ant. Idx	1	2	3	4
1	97.55	2.07	1.44	1.31
2	2.00	93.97	1.44	1.81
3	1.38	1.54	95.87	1.34
4	1.32	2.04	1.07	95.54

(a)

Ant. Idx	1	2	3	4
1	91.70	3.77	3.10	1.43
2	3.83	92.71	1.74	1.72
3	3.40	0.86	94.80	0.94
4	1.82	2.23	1.76	94.19

(b)

Figure 4: Confusion matrix between different TX antennas' channel patterns: (a) Using channel magnitude/phase distortion as antenna signature; (b) Using channel magnitude alone as antenna signature.

When both channel magnitude and phase are used as signatures, on average in 95.8% of cases, an antenna's instantaneous signature remains a best-match with its other signatures. Even with channel magnitude alone as signature, the matching probability remains around 93.3%. This clearly shows the feasibility and potential of antenna index modulation/decoding in Halma. Notably, since Halma aims to hop antennas on a symbol basis, even 1% confusion probability may result in decoding error across a packet and thwart any throughput gain. Thus, we must properly design the antenna hopping pattern to contain such errors.

3. Halma DESIGN

Halma consists of two key components: Antenna Index Coding (AIC) and Adaptive Antenna Hopping (AAH). AIC is a PHY module that creates an extra data stream by hopping through multiple antennas that share the same RF chain. AAH serves as a link-level module that adaptively picks the best set of antennas for AIC.

3.1 Antenna Index Coding (AIC)

3.1.1 AIC: an overview

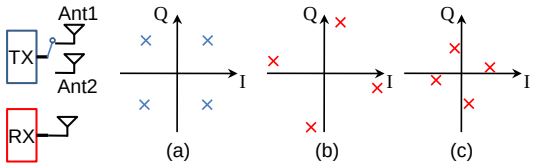

Figure 5: An example of AIC: (a) transmitted signal constellation; (b) received signal constellation for signals from antenna 1; (c) received signal constellation for signals from antenna 2.

At a high level, AIC works as follows. The transmitter divides its packet's data bits into two streams. The first stream is mapped to data symbols using a legacy modulation scheme, say QPSK. Before sending each symbol, the transmitter chooses which antenna to use for transmission, and the choice is driven by bits in the second stream, *e.g.*, '00' for antenna 0 and '10' for antenna 2. All data symbols of a packet are preceded by a *preamble* – a short sequence of known symbols, emitted sequentially by different antennas.

Upon detecting the packet, the receiver first extracts the "signature" of each transmit antenna based on the known preamble. The signatures differ, intuitively, in the way they distort the original QPSK constellation's phase/magnitude, as exemplified in Figure 5. These signatures then act as templates for decoding the antenna index hidden in the data symbols that follow. For each data symbol, the receiver can match the signal distortion pattern with the template signatures to decode the antenna index. Then, it normalizes the symbol by the signature, so that the symbol aligns with some point in the original QPSK constellation. Thereafter, the two bits in the symbol can be successfully decoded.

The above exposition abstracts out two non-trivial challenges in realizing AIC: *(i)* Real-world wireless systems do not modulate a data symbol as a single complex number. Instead, each symbol spreads over time (for single-carrier systems), or across frequency bins (for multi-carrier systems). While hopping between antennas, AIC must maintain integrity of the original symbols. *(ii)* Channel noise can cause variation of antenna signatures and thus decoding error when the receiver attempts to identify the antenna index. Erroneous antenna index may map the corresponding data symbol to a wrong constellation, thus triggering more bit errors.

Below we detail the design of AIC to meet the challenges.

3.1.2 Time-domain AIC for single-carrier systems

In single-carrier systems, each *data symbol* occupies the entire spectrum bandwidth, and its time-domain waveform comprises a sequence of 0-1 wavelets, called *chips*. Different data bits are mapped to *orthogonal* chip sequences. The receiver needs to decode individual chips, and then cross-correlates the resulting chip sequence with known sequences, the best match being remapped to desired data bits. This so called direct-sequence spread spectrum modulation (DSSS) scheme is used in the 802.15.4 ZigBee and the 802.11b WiFi standard. Without loss of generality, we design AIC on top of the ZigBee PHY-layer. Our design strikes a balance between AIC *efficiency* (number of hops per unit time) and *fault tolerance* (to channel noise and synchronization errors).

Sub-symbol level antenna hopping. To embed antenna index into data symbols, a straightforward approach is to switch antennas *per data symbol*, as in conventional SSK [5]. But this severely underutilizes AIC's potential, because a data symbol consists of multiple complex samples and, theoretically, it is possible to switch antenna *per sample* to convey more information per unit time.

Unfortunately, per-sample antenna hopping dramatically reduces the receiver's capability to decode the hidden antenna index, because ZigBee does not provide sample-level time/frequency synchronization. Even if synchronization can be achieved by upgrading the receiver hardware, channel noise can easily corrupt the antenna index.

AIC strikes a balance by using sub-symbol level antenna hopping. It forces the transmitter to use the same antenna for every N_s samples, where N_s falls between 1 chip (4 samples) and 1 symbol (32 chips). The receiver judiciously takes advantage of such redundancy across multiple samples to reduce antenna decoding errors.

A natural question here is how to configure the N_s. Ideally, N_s should be large enough to combat decoding errors through redundancy, yet small enough to harness the benefits of sub-symbol antenna hopping. Suppose the channel signatures of different transmit antennas are Gaussian *i.i.d.* random variables (corrupted by channel noise). Then the resulting antenna decoding error rate can be approximated by manipulating the Q-function for Gaussian random variables:

$$E = 1 - (1 - Q(\sqrt{SNR}))^{N_s}$$
$$\leq 1 - (1 - 0.5e^{-\frac{SNR^2}{2}})^{N_s} \text{ (Chernoff Bound)} \quad (1)$$

This simplified model implies that under a given SNR, *the decoding error bound decreases exponentially with the antenna hopping period* N_s. The estimation is roughly consistent with our empirical tests in real channel environment (Section 5.1). In AIC, we choose $N_s = 8$ as default, which results in a decoding error of only around 10^{-4} (Sec. 5.1).

Given a negligible decoding error, we can analyze *AIC's asymptotic link capacity gain* as follows. Legacy ZigBee represents every 4-bit data symbol by 128 samples, resulting in 0.0312 data bits per sampling period. In AIC, antenna hopping occurs per N_s samples and each antenna index represents $\log_2(N_t)$ bits. Thus, AIC boosts link capacity to $0.0312 + \log_2(N_t)/N_s$ bits per sample. Consequently, AIC achieves a capacity gain of $(0.0312 + \log_2(N_t)/N_s)/0.0312 = 1 + 32\log_2(N_t)/N_s$ over ZigBee. For instance, with only two transmit antennas ($N_t = 2$) and $N_s = 8$, theoretical capacity gain over legacy ZigBee can be $5\times$.

Antenna index decoding. To decode the transmit antenna index, ideally, the receiver should estimate both the channel magnitude and phase distortion *w.r.t.* each transmit antenna. But this would require substantial hardware modification to ZigBee, whose PHY uses a non-coherent demodulation scheme and simple correlation-based decoder that requires no channel estimation. Our AIC decoder circumvents this issue using a *template matching* mechanism, leveraging an inherent structure of ZigBee packets.

We observe that the modulated waveform of any chip sequence is made from 4 elementary patterns, corresponding to 32 complex samples (Figure 6 shows two of such patterns). Thus, each transmit antenna only needs to send these 32 samples as a *training template*, piggy-backed in the beginning of each packet (Figure 6), to facilitate the receiver's decoding.

The receiver first decodes normal ZigBee data symbols across a packet, and then reverts to the beginning of the packet to decode the antenna index embedded in every N_s

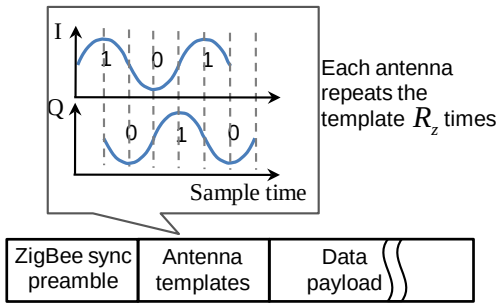

Figure 6: Packet format and time-domain antenna hopping for ZigBee.

samples. Note that, by this time the receiver already knows which symbol each group of samples represent. Thus, it can match the samples with those template samples representing the same symbol, but sent by different antennas. The antenna with the most similar template is most likely used by the sample.

To combat signal variations caused by channel noise, each transmit antenna repeats each training template R_z times, and the receiver uses the average of the R_z repetitions as one template. Further, the receiver harnesses AIC's embedded redundancy – it runs a majority vote among the estimated antenna indices of the N_s samples, so as to determine the transmit antenna they use.

AIC uses Euclidean distance between sample values as a metric for template matching. More formally, lets denote T_{i,s_j} as the antenna template from transmit antenna i while sending a symbol s_j. Each raw sample y_j within the N_s-sized group, has already been decoded as sample s_j. Then the antenna index decoding is represented by:

$$I = Mode_{s_j \in \mathbf{N_s}} \{ \arg \min_{\forall i \in \mathbf{N_t}} ||y_j|^2 - |T_{i,s_j}|^2| \} \qquad (2)$$

where $Mode\{\cdot\}$ denotes majority vote over a set, and N_t is the number of TX antennas. Notably here, decoding the antenna index only requires $N_s N_t$ operations in total.

The following points are worth noting for AIC decoding:

(i) Why decoding normal data symbols and antenna indices separately? Since ZigBee uses a differential demodulator to decode normal data symbols, and antenna switching occurs only per $N_s = 8$ samples, AIC itself is unlikely to affect the performance of the normal decoder. Therefore, the receiver decodes the normal data symbols first, *separately* from antenna index decoding.

(ii) Preamble overhead. Compared with ZigBee, the only overhead lies in the antenna templates, sent sequentially by N_t antennas and repeated R_z times. Each antenna template only contains 32 complex samples. With $N_t = 4$ and a default setting $R_z = 4$, the total overhead is 512 complex samples — equivalent to only 128 μs, and less than half of ZigBee's legacy preamble length.

Tolerating lack of synchronization. The above description implicitly assumed the receiver knows the exact position of each data sample/chip. In reality, the ZigBee packet preamble only ensures coarse, symbol-level synchronization. Sampling time offset between the transmitter and receiver does not significantly affect the ZigBee decoder that uses correlation based decoding, yet it can cause smearing of adjacent samples, thus increasing the antenna error rate (AER). Sampling-offset compensation is possible but will increase the receiver complexity. The carrier frequency synchronization between transmitter and receiver bears a similar issue.

Halma has two inherent counter-measures to the lack of fine-grained sampling time/frequency synchronization. First, after grouping N_s samples and performing a majority vote, impact of the sampling offset is reduced as it only affects samples near the boundary of the group. Second, legacy ZigBee repeats a known chip sequence 6 times in its preamble. The receiver achieves coarse synchronization by finding the *first* sequence using correlation. In Halma, the receiver uses a sliding-window based correlation for all 6 chip sequences in the preamble. It finds the correlation position that minimizes the maximum number of chip errors, and uses that position as a sync point. This simple extension can synchronize transmitter and receiver within one chip, equivalent to two raw complex samples.

3.1.3 *Frequency-domain AIC for multi-carrier systems*

Encoding antenna index across subcarriers. In WiFi OFDM systems, bits are first modulated into data symbols following certain constellation, *e.g.*, QPSK. Then, each data symbol, represented by a complex sample, is modulated onto a frequency bin called subcarrier. A group of subcarriers forms an OFDM symbol, and a group of OFDM symbols forms a packet. Note that the L data symbols (*e.g.*, 48 for 802.11g, 52 for 802.11n) embedded in an OFDM symbol are inseparable in time domain, yet antennas can only be switched over time. This dilemma inspires us to migrate AIC to the frequency domain.

Specifically, we assign different subcarriers to different antennas to emulate antenna switching in frequency domain. Figure 7 illustrates an example with 2 TX antennas. Each subcarrier can be occupied by only one transmit antenna, and index of that antenna conveys extra bits of information. Similar to the time-domain AIC, we maintain robustness by forcing N_f adjacent subcarriers to share the same antenna. Said differently, antenna switch occurs only for every N_f subcarriers. N_f has to be a divisor of L and is default to 6.

Notably, all subcarriers in an OFDM symbol fully overlap with each other in time, and therefore, all transmit antennas need to be active simultaneously. In other words, frequency-domain AIC requires multiple RF chains at the transmitter, although the receiver is still single-antenna, single RF chain. From energy efficiency perspective, it will be most applicable for infrastructure wireless LANs, which are dominated by downlink traffic [9]. With Halma, single-antenna clients can benefit from throughput gain without costing extra energy or hardware.

Decoding frequency-domain antenna index.

(i) Synchronization and channel estimation. We leverage the built-in 802.11n packet preambles for synchronization and channel estimation (Figure 7). Specifically, an STF (short-training field) preamble, with periodic patterns in time-domain, is used for the receiver to detect the start of a packet [9]. An LTF (long-training field) preamble, with a known random sequence repeated twice, is used to first estimate frequency offset, and then estimate per-subcarrier channel gain (magnitude/phase distortion). Right after STF, each transmit antenna sends the LTF sequentially, and their channel estimation is used as antenna signatures at the receiver.

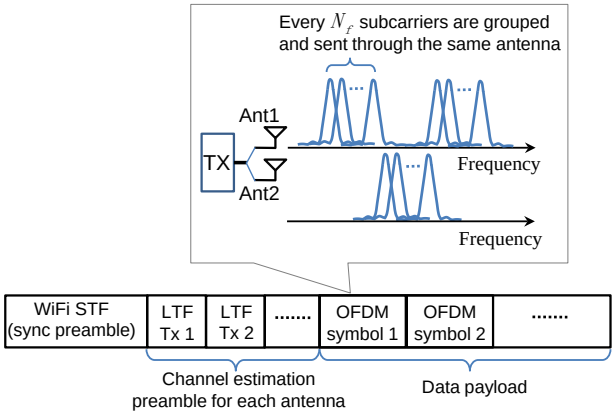

Figure 7: Antenna index coding (AIC) for OFDM WiFi. Antenna hopping occurs in frequency domain.

(ii) Joint decoding of data symbol and antenna index. Unlike the time-domain AIC, now the receiver has full synchronization, and both channel magnitude and phase pattern with respect to each TX antenna, which together enriches the TX antennas' signature space. Accordingly, the decoder takes advantage of this during antenna index decoding.

Denote $h_{i,l}$ as the channel gain for subcarrier l for TX antenna i. Let s_j be the j-th modulated data symbol, and y_l the received symbol in subcarrier l. Then, the receiver decodes the antenna index by finding the index that gives the minimum Euclidean distance, among all N_f subcarriers sharing the same antenna, *i.e.*,

$$I = \arg\min_{\forall i}(\sum_{l=k}^{k+N_f-1} \min_{\forall j} |y_l - h_{i,l}s_j|^2) \quad (3)$$

where $k = \{1, N_f+1, \ldots, L-N_f+1\}$ is the index of first data symbol within the subcarrier group sharing the same antenna. i and j index the TX antennas and modulated data symbols, respectively. After decoding the antenna index I, we use its channel information to normalize all N_f subcarriers inside the group, map the resulting symbol to its constellation, and decode the bits therein. Because of the channel gain normalization, the antenna index decoding and data symbol decoding are coupled. Therefore, WiFi is more sensitive to antenna decoding errors than ZigBee.

(iii) Overhead and asymptotic capacity gain. To facilitate channel estimation, N_t LTF preambles are needed per packet, equivalent to $8N_t$ μs overhead. This is negligible compared with a typical packet duration, even with $N_t = 8$.

For legacy WiFi, each sample represents M data bits under a modulation order of M. AIC can augment an additional $\log_2(N_t)/N_f$ bits per subcarrier on top. With $N_t = 4, N_f = 6$ and BPSK modulation ($M = 2$), the capacity gain is 33%. But with 64-QAM, the gain reduces to 6%. Thus, higher-order modulation in WiFi may marginalize the gain from AIC. In practice, however, WiFi links do not always utilize the optimal modulation. For example, signaling packets (ACK, RTS/CTS, *etc.*) are typically sent using BPSK, leaving sufficient link margin for Halma to establish an additional "control channel" [13] through antenna hopping. In addition, even under the optimal optimal modulation order, Halma can still harvest non-trivial throughput gain through its link-level adaptive antenna hopping, as explained below and verified in Section 5.2.

3.2 Adaptive Antenna Hopping (AAH)

For AIC to achieve high decoding confidence, the TX antennas' signatures should be as "dissimilar" as possible. However, if two antennas with highly disparate channel gains are used, the one with relatively low magnitude and hence low SNR, may bottleneck the system throughput. In order to strike a balance between channel dissimilarity and quality, the transmitter employs AAH to strategically hop between the optimal subset of antennas.

3.2.1 Adaptation protocol

The adaptation protocol in AAH consists of 3 key steps. Without loss of generality, we describe it for WiFi only.

(i) In the very beginning of AAH, the transmitter sends a polling packet with all N_t antennas sequentially sending LTF, the channel-estimation preamble.

(ii) A WiFi receiver extracts the channel gain (magnitude/phase) and noise level from the LTF. Then it estimates an optimal antenna *configuration* across three dimensions (antenna combination, number of subcarriers per antenna symbol N_f, and modulation size M) to maximize throughput. The receiver then informs the transmitter to use this configuration in subsequent AIC transmissions.

(iii) The optimal configuration may vary due to channel variation. Thus, the receiver monitors the throughput $TH(t)$ for current configuration. If its deviation to the initial throughput $|\frac{TH(t)}{TH_0} - 1|$ is larger than a certain threshold σ (we use an empirical value 0.1 by default), then the configuration is outdated, and the receiver requests the transmitter to resend the polling packet as in (i).

The above only sketches the basic AAH operations. The major challenge lies in step *(ii)*, which requires predicting the performance of a given configuration, and searching for the optimal configuration. We address these two problems through a model-driven framework, described below.

3.2.2 Modeling the AER and BER of AIC

In AIC, there are two signal spaces, the antenna index space Ω_a and the symbol space Ω_s. Let ε_a and ε_s denote the antenna error rate (AER) and bit error rate (BER) of these two signal spaces. The overall bit error rate can be calculated as $\varepsilon = \frac{\varepsilon_s\mu_s + \varepsilon_a\mu_a}{\mu_s+\mu_a}$, in which μ_s and μ_a are the number of bits for each symbol in the two signal spaces. For OFDM AIC, $\mu_s = log_2(M)$ and $\mu_a = log_2(N_t)/N_f$. Recall M is the modulation size, N_t the number of TX antennas, and every N_f subcarriers use the same antenna.

We define the SNR of antenna index decoding as,

$$SNR^a_{i,l,j} = d^2_{i,l,j}N_f/N_0, \quad (4)$$

where N_0 is the variance of receiver noise, modeled as zero-mean complex Gaussian noise. N_0 can be estimated from the received LTF as [12]:

$$N_0 = \sum_{i=1}^{N_t}\sum_{l=1}^{L} |\widehat{LTF^1_{i,l}} - \widehat{LTF^2_{i,l}}|^2/(LN_t) \quad (5)$$

in which $\widehat{LTF^1_{i,l}}$ and $\widehat{LTF^2_{i,l}}$ are the LTF symbols of subcarrier l and antenna i, in the 1^{st} and 2^{nd} half of the LTF respectively, which are identically modulated. L is the number of data subcarriers in one OFDM symbol.

$d_{i,l,j}$ is the minimum Euclidean distance between channel gain of antenna i and other antennas on subcarrier l for symbol j, *i.e.*,

$$d_{i,l,j} = \min_{\forall m\neq i,n} |h_{i,l}s_j - h_{m,l}s_n|, \quad (6)$$

where $h_{i,l}$ is the channel gain of subcarrier l in TX antenna i. s_j is the modulated symbol.

Given the antenna index decoding SNR in Eq. (4), the AER is simply the probability that one Gaussian random variable smears into the other's "region", which can be modeled by the standard Q-function. Consequently, we can model AER as:

$$\varepsilon_a = \frac{\sum_{i=1}^{N_t}\sum_{l=1}^{L}\sum_{j=1}^{M} Q(\sqrt{SNR^a_{i,l,j}})}{N_t L M},$$

where we average AER over L subcarriers and N_t antennas.

For the BER, we adopt the same model as the effective SNR, which has proven to be accurate [14]. Details are omitted to avoid duplication.

As for ZigBee, the above model can be applied with few minor modifications. Specifically, as ZigBee only employs channel magnitude to decode the antenna index, the antenna index decoding SNR is defined as,

$$SNR^a_{i,j} = d^2_{i,j} \cdot (N_0|\mathbf{S}|)^{-1} \quad (7)$$

Here $\mathbf{S}$ denotes the set of samples transmitted through the "antenna templates". The distances $d_{i,j}$ is defined by the minimum Euclidean distance between the corresponding antenna templates from antenna i to other antennas, for all template O-QPSK symbols:

$$d_{i,j} = \min_{m\neq i, \forall m\in\mathbf{N_t}, \forall s_j\in\mathbf{S}} ||T_{i,s_j}|^2 - |T_{m,s_j}|^2| \quad (8)$$

where T_{i,s_j} denotes the "antenna template" as defined in section 3.1.2.

The noise N_0 is calculated using the variance of all template symbols for an antenna as:

$$N_0 = \sum_{\forall i\in\mathbf{N_t}, \forall s_j\in\mathbf{S}} Var(T_{i,s_j}) \quad (9)$$

The AER is then represented as,

$$\varepsilon_a = (N_t|\mathbf{S}|)^{-1} \cdot \sum_{i=1}^{N_t}\sum_{j=1}^{|\mathbf{S}|} Q(\sqrt{SNR^a_{i,j}}) \quad (10)$$

ZigBee AIC isolates antenna decoding from normal data demodulation (Section 3.1.2), we thus only need to model AER for the AAH protocol.

3.2.3 Model-driven adaptation algorithm

Based on the model in Section 3.2.2, a receiver can map the overall bit error rate ε to expected throughput under a given configuration. For simplicity, we assume no error correction code is adopted. Then, the packet level throughput can be modeled as:

$$Th = \frac{PacketSize * (1-\varepsilon)^{PacketSize}}{PacketDuration} \quad (11)$$

To obtain the throughput-optimal configuration from the model, one approach is to search all possible configurations. But this results in a formidable computational complexity of $O(2^{N_t}N_t^2M^2|\mathbf{N_f}|)$, where $|\mathbf{N_f}|$ is the cardinality of the set of possible antenna switch rate.

Therefore, we design an efficient algorithm that approaches the best configuration in two tractable steps, aiming to strike a balance between channel quality and dissimilarity.

First, we generate a series of combinations of antennas $\mathbf{C} = [\mathbf{N_1}, \mathbf{N_2}, \mathbf{N_4}, \ldots, \mathbf{N_{2^i}}, \ldots]$, where $\mathbf{N_{2^i}}$ denotes a combination of 2^i antennas, and $i = 0, 1, \ldots, \lfloor \log_2(N_t) \rfloor$. These 2^i antennas are chosen with the highest effective SNR out of all possible antennas. In the second step, we estimate the corresponding throughput for the antenna combinations in set C with different modulation size M and switch frequency $1/N_f$. The configuration that gives the highest throughput will be conveyed to the transmitter.

These two steps essentially constitute a greedy Algorithm driven by the throughput-model, which we summarize in Algorithm 1. The algorithm reduces the antenna searching space from 2^{N_t} to $\log_2(N_t)$, resulting in an overall computation complexity of $O(\log_2(N_t)N_t^2M^2|\mathbf{N_f}|)$. Since N_t and M are usually small (below 8), the computation cost is negligible under practical settings and with our empirical configuration of $N_f = 6$. ZigBee's AAH protocol follows similar mechanisms, and is omitted due to space constraint.

Algorithm 1 Model-driven Adaptation Algorithm for AAH

```
Receive ⟨CSI, N0⟩
foreach antenna i
    Compute effective SNR SNR_i^s
end foreach
Generate C
Max_TH = 0
foreach N_2^i in C
    Calculate ⟨Th, Nf, M⟩ such that,
    Th = max_{∀Nf,M} Th(ε)
if Th > Max_TH
    ⟨Max_TH, Best_ant, Best_Nf, Best_M⟩ =
    ⟨Th, N_2^i, Nf, M⟩
end if
end foreach
Return ⟨Best_ant, Best_Nf, Best_M⟩
```

4. IMPLEMENTING Halma

We have prototyped Halma's AIC and AAH modules on the WARP software radio platform [15]. Our implementation realizes both the single-carrier Halma for ZigBee and multi-carrier for WiFi.

Halma for ZigBee Transceiver. We port an open-source C++ implementation of ZigBee PHY layer [16] to the WARPLab driver. This implementation is validated by running it on WARP and allowing direct communication with a COTS ZigBee transceiver [8]. On top of it, we develop the single-carrier AIC and its decoding mechanisms following Section 3.1.2, along with the AAH protocol (Section 3.2.1).

Halma for WiFi Transceiver. To verify Halma for multi-carrier systems, we first implemented an 802.11n-compliant OFDM communication library consisting of a (i) transmitter module: bit-to-symbol mapping (PSK/QAM), OFDM modulation, preamble/pilot embedding; (ii) receiver module: packet detection, synchronization, frequency offset compensation, and OFDM/symbol demodulation functions. The multi-carrier AIC encoding/decoding and AAH modules (Section 3.1.3 and 3.2.1) are then implemented on top of the 802.11 PHY library. Our implementation reuses the 802.11n MIMO preamble mechanism, that allows transmit antennas to send LTF preambles sequentially, from which the receive antenna can extract their channel pattern.

As a benchmark comparison, we have also implemented the 802.11n STBC scheme that exploits diversity gain between a multi-RF-chain transmitter and single RF-chain receiver. We further integrate Halma with STBC, by allowing

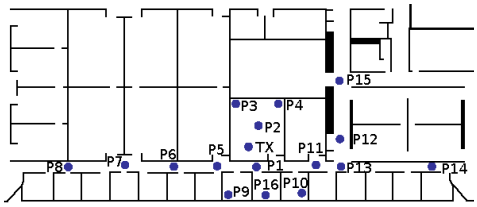

Figure 8: Testbed topology.

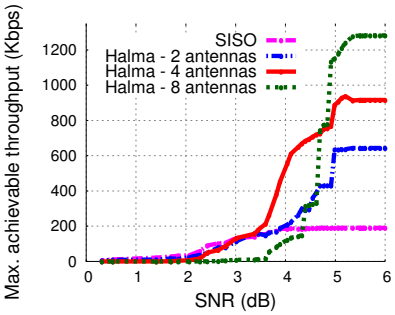

Figure 9: Throughput of ZigBee using SISO and Halma.

the transmitter to hop between different pairs of transmit antennas, each pair running the STBC based modulation.

5. EVALUATION

We evaluate the effectiveness of Halma's design in a testbed with 6 WARP boards, each having 4 antennas. Two of the boards can form an 8-antenna transceiver using WARP's clock expansion module. The testbed is configured to an unused WiFi channel 14 to isolate ambient interference. Figure 8 shows the floor plan of our testbed, where nodes are moved around 16 different locations to create a larger topology.

5.1 Performance of Halma for ZigBee

Raw throughput performance. We begin with a micro-benchmark throughput comparison between Halma and ZigBee (SISO). The throughput metric here computes the net throughput after the impact of Halma's preamble overhead and packet losses caused by AER or BER. Packet size is configured to its maximum (128 bytes). Channel condition is gauged by the receiver according to its Link Quality Indicator (LQI), which can be calculated based on chip error rate and converted to SNR following the mapping table of TI CC2420 [17].

Figure 9 plots the throughput under different SNR conditions created by varying link distance. We disable the AAH mechanism, and randomly select a set of transmit antennas to run Halma. But for SISO, we use an *exhaustive search* to pick the antenna resulting in highest throughput, in consistent with antenna selection mechanisms in legacy devices [8]. We observe that under ultra-low SNR (below 3.5 dB), Halma's throughput is comparable to ZigBee, or even lower due to high antenna error rate (AER). However, in the common SNR range above 5 dB, it achieves $3.1\times, 4.7\times$ and $6.4\times$ throughput gain, with 2, 4, and 8 antennas, respectively. Owing to Halma's sub-symbol level AIC, the gain can be super-linear for 2 and 4 antennas, consistent with our analysis in Section 3.1.2. With 8 antennas, AER becomes non-trivial because of similarity in antenna signatures, and thus a sub-linear gain is achieved.

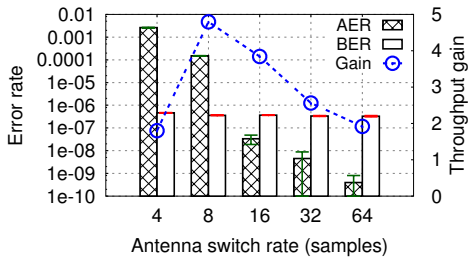

Figure 10: Impact of antenna switching granularity. Number of antennas $N_t = 4$. SNR >5 dB.

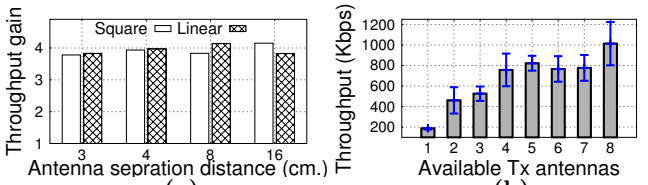

Figure 11: (a) Impact of transmitter's antenna separation and placement patterns. $N_t = 4$. (b) Throughput as a function of available number of TX antennas.

Granularity of antenna hopping. As mentioned in Section 3.1.2, fine-grained antenna hopping, ideally on a per-sample basis, can deliver more bits per unit-time. Yet it exacerbates the AER. This tradeoff is manifested in Figure 10. A sweet-spot of 8-samples exists and is used as the switching granularity across our evaluation. In addition, BER is virtually unaffected by the switching granularity, while AER decreases exponentially as we increase the antenna-switching period. Both observations are consistent with the premise behind Halma's AIC design.

Effect of transmit antenna separation and placement. In this experiment we try to identify the antenna separation needed to create distinct signatures for Halma to work. We place the receiver 1 m away from the transmitter within line-of-sight, and vary the transmitter's antenna separation. Figure 11(a) shows that a small separation of 3 cm is sufficient to create a significant difference in the channels so that the receiver can discern the antenna index with high probability, thus achieving high throughput gain over ZigBee SISO. Also, different antenna placement patterns do not noticeably affect the performance. This implies that Halma can be deployed on a 4-antenna device with an area of 3cm×3cm, which may be suitable for small sensor nodes.

Throughput gain *w.r.t.* number of available antennas. When the transmitter has the luxury of using a larger number antennas, it owns more options to select the best group of antennas. In this experiment, we vary the number of available antennas at the transmitter side and run an exhaustive search to pick the best antenna group that achieves the highest throughput. The receiver is placed at 1m away from the transmitter. Figure 11(b) shows that throughput increases as the number of transmit antennas increases. Yet the trend may saturate, primarily because it becomes harder to ensure dissimilarity between antennas.

Energy efficiency of Halma. We now evaluate the energy efficiency of Halma in ZigBee through trace-driven simulation. We first simulate a ZigBee WPAN cell containing 20 nodes under ZigBee's TDMA mode. The packet arrival time of each nodes is modeled as a Poisson process with mean arrival time 0.050s (=1/20). The TDMA schedule bears a

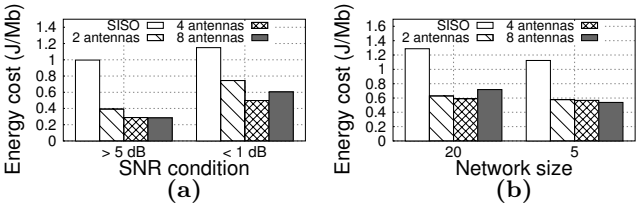

Figure 12: Energy cost in ZigBee: (a) TDMA. (b) CSMA.

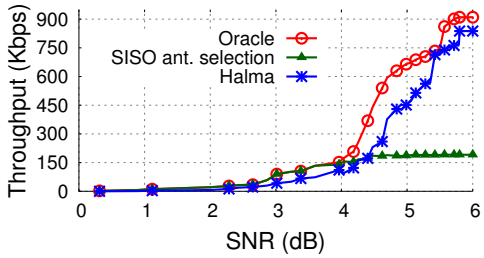

Figure 13: Performance of Halma's AAH protocol compared with SISO antenna selection and oracle.

wake-up interval of 0.04s to support the packet transmission. We inject the power measurements of the multi-antenna ZigBee board (Section 2), along with the throughput statistics in Figure 9, into the simulator to evaluate the total energy consumption. The results (Figure 12(a)) show that in common SNR conditions (all links >5 dB), Halma's energy cost is 60% lower than SISO even with 2 antennas, which is mainly attributed to the much shorter time spent in transmission. In poor SNR condition (< 1 dB), Halma's energy cost increases but can still save more than 29% energy for SISO.

In another experiment we fix the SNR to >4 dB, and simulate ZigBee's CSMA MAC protocol under two different network sizes. Notably, for for both network sizes, 2-antenna Halma can save around 50% of energy compared with SISO. Yet improvement from 2-antennas to 4-antennas is insignificant (Figure 12(b)). This is mainly because the idle listening energy consumption becomes non-trivial in CSMA, which partly nullifies Halma's throughput gain — although more antennas allow Halma to finish transmission quickly, the idle time also increases.

Effectiveness of AAH. To verify the AAH design for ZigBee Halma, we move the receiver to different locations to create a variety of SNR conditions. The experiments are conducted in an busy office environment with 12 people, and 2 intentionally walking back and forth. We compare AAH with an Oracle scheme that searches the best set of antennas offline (based on packet traces). Figure 13 shows that AAH's model-driven algorithm can closely track the Oracle for different SNR conditions. We also ran the RSSI-based antenna selection for SISO as described in [18]. Since this approach only picks a single antenna without AIC, Halma outperforms it by 1.9× to 4.4× in common SNR ranges.

Cumulative gain. Figure 14 plots the CDF of Halma's throughput gain over legacy ZigBee (which uses RSSI-based antenna selection [18]), across all receiver locations in our testbed map. Halma delivers significant throughput gain for majority of the locations. With 2 TX antennas, it outperforms SISO in > 80% locations, and achieves more than 1.5× gain for more than 60% locations. With 4 antennas, the gain reaches 3× for more than 89% of the case. The

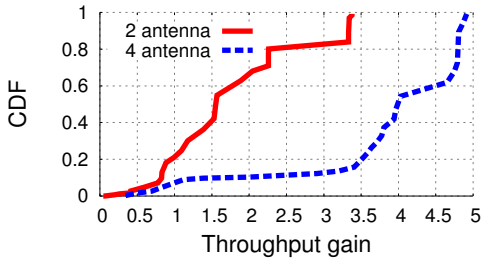

Figure 14: CDF of throughput gain over SISO (ZigBee).

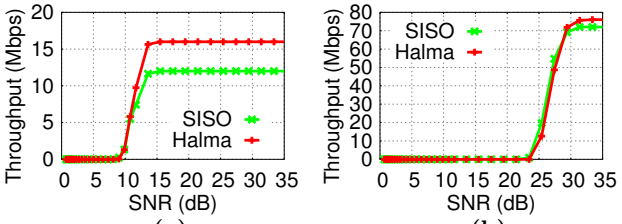

Figure 15: Achievable throughput for a fixed modulation size. (a) M = 2. (b) M = 64.

remaining small fraction of cases encounter high antenna similarity, thus even lower throughput than SISO.

5.2 Performance of Halma for WiFi

Achievable throughput. We evaluate the throughput of different WiFi modes: SISO (with RSSI-based antenna selection [18]), STBC, Halma and Halma-STBC. By default, Halma uses a given set of $N_t = 4$ antennas, antenna switching rate $N_f = 6$ subcarriers and packet size 1 KB. AAH is disabled in this experiment. As we can see from Figure 15, in the common SNR region where packet loss rate is low and throughput stablizes, Halma can achieve around 32% throughput gain when running over BPSK. With higher-order modulation like 64-QAM, Halma's gain is marginalized, but it still adds an extra 4 Mbps to SISO, which is sufficient to create a free control channel as in Flashback [13].

Impact of antenna switching frequency. Figure 16 shows the impact of antenna switch frequency N_f. We see a similar tradeoff between switching rate and AER as in ZigBee Halma. A very high antenna switching rate (*e.g.*, 2 subcarriers per switch) results in higher antenna decoding error rate, and hence drastically lowers throughput. Notably, high AER may trigger larger BER as incorrect channel distortion is compensated in the symbol decoding of WiFi. When $N_f \geq 6$, the frequency-domain AIC achieves high decoding confidence, with AER comparable or even lower than OFDM BER. The phenomenon is unaffected by modulation rates.

Throughput vs. SNR. We further evaluate the achievable throughput when modulation rate adaptation is enabled. The experiments run over channel traces collected under a variety of SNR conditions, and the best modulation rate is computed for both SISO and Halma offline. The result (Figure 17(a)) shows that such an ideal rate adaptation scheme marginalizes the gain from Halma if its AAH is disabled. But even in this case, Halma can exploit the link margin between modulation levels [13] to deliver several extra Mbps of throughput. When AAH is enabled (Figure

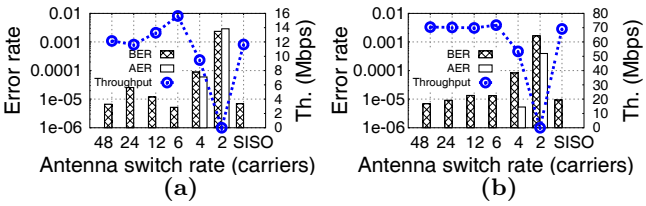

Figure 16: AER and BER vs antenna switch frequency: (a) M = 2. (b) M = 64.

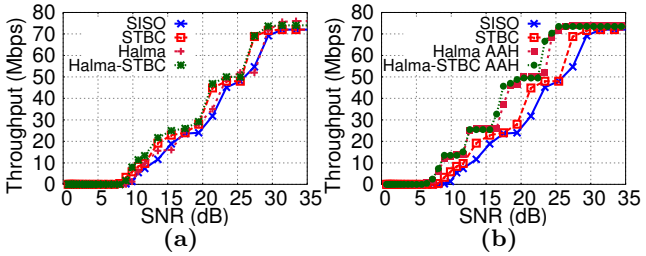

Figure 17: System throughput vs. SNR, with modulation rate adaptation.

17(b)), it delivers up to 90% of throughput gain and more than 30% in most SNR conditions. Notably, STBC does not add significant improvement to either Halma or SISO, mainly because it is used for combating small-scale fading, which mostly manifests in high-mobility scenarios.

Accuracy of throughput model. Recall the AAH module adopts a model-driven approach to predict achievable throughput (Section 3.2.2). We evaluate the model by comparing it with an oracle that computes the maximum throughput offline by searching across all modulation sizes, N_f and antenna groups (among 4 antennas). Figure 18 shows that the throughput model closely approximates the oracle, and thus it can be instrumentally used for the AAH protocol. Notably, *if AAH is not used and all antennas are greedily selected (labeled as "All Antennas"), then the performance can be degraded by a median value of 45%.* This again substantiates the importance of balancing link quality and channel dissimilarity, which has not been exploited in prior work.

Impact of the number of antennas. In this experiment, we vary the number of available antennas to evaluate the achievable throughput of the system. As a microbenchmark evaluation, the achievable throughput is evaluated by first collecting channel traces, and then exhaustively searching over all possible set of antennas offline. Figure 19 shows that, with 2 and 4 antennas, Halma achieves 38% and 60%

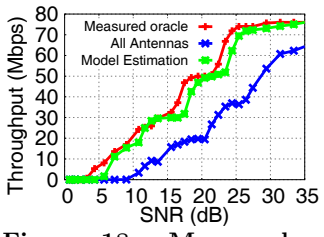

Figure 18: Measured oracle throughput and the modeled throughput over different SNR.

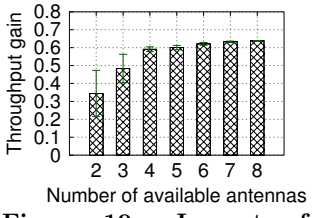

Figure 19: Impact of the number of available antennas on Halma's throughput.

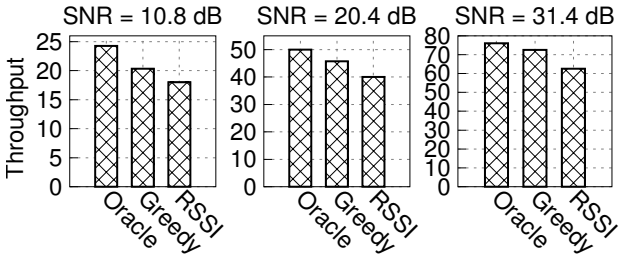

Figure 20: Measured throughput of Halma's AAH, Oracle and RSSI-based SISO antenna selection.

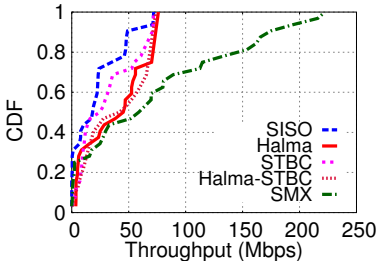

Figure 21: Cumulative distribution of the throughput for different schemes.

throughput gain over WiFi SISO. Similar to ZigBee-Halma, further increasing the number of antennas brings marginal gain, partly because of the increasing antenna similarity, and partly because Halma's AAH only picks antennas with highest modulation rates in WiFi.

Effectiveness of AAH for WiFi. Similar to the ZigBee setting, Figure 20 compares Halma's greedy AAH adaptation protocol with the Oracle and RSSI-based SISO antenna selection mechanism [18]. Greedy achieves 80% to 97% of the Oracle throughput across different SNR levels and 1.13× to 1.21× higher than the SISO antenna selection.

Performance in the field test. We conduct an integrated test of Halma for WiFi in a dynamic office environment similar to the ZigBee setting. Figure 21 plots the resulting throughput distribution of all testbed locations. We see that Halma outperforms SISO with antenna selection (1.45× on average), and even STBC, for almost all the locations. Its mean throughput is lower than MIMO spatial multiplexing (SMX), which utilizes 4 antennas at the receiver. Notably, for low-SNR users, Halma can have comparable performance to SMX, because it adaptively chooses the high-link-quality antennas rather than being bottlenecked by the antenna with low gain. In addition, although SMX runs 4 RF chains concurrently (approximately 3× energy cost, see Section 2), its mean throughput gain over Halma is only 1.4×. This implies that SMX's energy-per-bit is still much higher than Halma under saturated traffic conditions.

Energy consumption under WiFi workload. To estimate the energy consumption of Halma under practical traffic patterns, we use a trace-driven approach similar to [9]. The WiFi packet traces are collected from (i) an FTP session downloading a 25 MB file, (ii) a 5-minute web browsing session, (iii) a 5-minute VoIP session using Google+ hangout. We replay the traces using the power statistics of the Atheros 9380 card, along with the bit-rate statistics collected from our throughput experiments (with an intermediate SNR of 20.4 dB). The bit-rate of SMX is 3× of SISO (Figure 21). From the results (Figure 22), we make two observations.

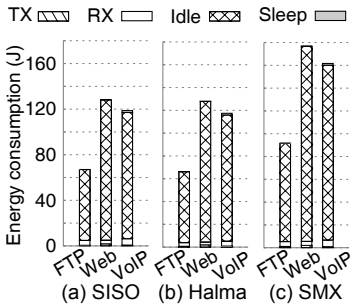

Figure 22: WiFi energy consumption for different schemes under the same amount of traffic load.

First, Halma consumes comparable energy as SISO, despite its higher throughput shown in prior experiments. Thus, the throughput gain of Halma in WiFi does not translate into energy saving (unlike in ZigBee), primarily because only 10% of channel time is spent in idle listening under practical WiFi traffic patterns [9]. Second, MIMO SMX consumes much higher energy than Halma and even SISO, despite its much higher throughput. Since the *same amount of traffic* is delivered by different schemes, this result implies that Halma is much more energy efficient than MIMO SMX under realistic WiFi traffic patterns.

6. DISCUSSION

Where is Halma applicable? From the foregoing experimental evaluation, we conclude that for single-carrier communication devices like ZigBee sensors, Halma can substantially improve link throughput. This improvement can be directly translated into energy reduction since both transmitter and receiver maintain a single RF-chain. Admittedly, Halma requires multiple antenna elements at the transmitter side, entailing more space cost. However, we have observed vast energy saving from Halma even with 2 transmit antenna elements, with marginal space cost. Such multi-antenna-element ZigBee nodes already exist [8]. On the other hand, for multi-carrier communications devices like WiFi, Halma requires multiple RF chains at the transmitter side to achieve throughput gain. However, since WiFi is dominated by downlink traffic originating from the energy-insensitive access point, Halma's throughput gain can still benefit single-antenna clients without increasing their energy cost.

Higher-order modulation for ZigBee? As we analyzed in Section 3.1.2, higher order modulation schemes like 64-QAM may scale link capacity just like Halma. Thus, one may wonder why ZigBee hardware does not support such modulation levels. The reason again lies in energy cost. Higher-order modulation schemes intentionally vary both data symbol amplitude and phase to convey information, which requires power-hungry linear amplifiers in the RF chain [19, Ch. E3]. Low-level modulation, including BPSK and the default O-QPSK in ZigBee, manifests a constant-envelop waveform, thus enabling simple and highly efficient non-linear amplifiers [19]. In some sense, Halma actually augments amplitude modulation on legacy ZigBee by leveraging the symbol amplitude variation naturally provided by the wireless channel — the channels between different TX antennas and the RX antenna. Thus, it does not need the costly power amplifier.

Antenna switching overhead. Antenna switch has already been equipped on many WiFi and ZigBee devices [8], although it is mainly used to select antennas on a coarse-grained manner (every a few packets). Commercial-Off-The-Shelf antenna switches typically consume several μW of power – orders of magnitudes lower than TX/RX/idle power [6, 20]. Their response time falls within a few ns – negligible compared with the switching period in Halma (8 samples or 4 μs for ZigBee). Therefore, Halma's fine-grained, sub-symbol-level antenna switching mechanism is feasible in practice. In fact, the Atmel multi-antenna ZigBee receiver [8] uses an antenna switch to decide which antenna to use immediately after a packet preamble is detected. The switching latency is negligible and completely hidden from the ZigBee demodulator. Note that a WiFi transmitter running Halma still needs multiple active antennas (Section 3.1.3), and the antenna switch is used only by AAH on a per-packet basis.

7. RELATED WORK

Communication by switching antennas. Halma is partly inspired by the communication-theoretic concept of Space-Shift-Keying (SSK) [5,21,22], also referred to as Spatial Modulation (SMod) when augmented on top of narrowband PSK/QAM modulation mechanisms [23]. A solid theoretical foundation has been established that justifies the potential capacity gain of SMod over SISO (See [24] for a theoretical analysis, [6] for a comprehensive survey and [25] for a first measurement validation). We have thoroughly discussed Halma's unique advantages over conventional SSK (Sec. 1), particularly in its asymptotic gain in wide-band single-carrier and multi-carrier systems. To our knowledge, Halma is the first scheme that reveals these observations in a real implementation and unleashes the potential of antenna hopping for single RF-chain transceivers.

Communication through side channels. Besides traditional modulation schemes, recent wireless networks witnessed many novel cross-layer communications schemes that exploit side channels. 802.11ec [26] employs short, correlatable symbol sequences to replace RTS/CTS, thus reducing the control message overhead. Flashback [13] embeds high-power single-tone signals into OFDM subcarriers, so as to create an extra control channel (with up to 400Kbps rate) on top of the normal data transmission. SideChannel [27] allows a transmitter to modulate energy pulses on top of an existing transmitter's packet, which can be identified by the receiver and improve ZigBee capacity by 2.5×. Both Flashback and SideChannel exploit the link margin between practical, conservative modulation protocols and an oracle choice. Similar to such schemes, the bonus bit-rate resulting from Halma's antenna index modulation can be applied to create a covert channel. Owing to multiple antennas, Halma's bonus channel demonstrats a much higher capacity.

Antenna selection for MIMO networks. Halma's adaptive antenna hopping protocol inherits the insights from MIMO antenna selection. Information theoretic analysis has predicted the asymptotic SNR improvement from antenna selection to be $\log(N_t)$ times [28,29], assuming i.i.d. channel fading. Practical antenna selection protocols [18,30] tend to pick a single best antenna based on link quality estimation. In Halma, a transmitter adaptively picks a set of antenna to hop between, using a model-driven approach. Combined with antenna index modulation, it achieves much higher net-

work throughput compared with traditional antenna selection schemes (Section 5).

MIMO link energy optimization. Many MAC-layer protocols [3, 4, 10, 31] have been proposed that adaptively choose the number of RF chains to balance the throughput and energy consumption of WiFi MIMO transceivers. Halma sticks to a single RF-chain receiver, and consumes similar energy as SISO under real traffic patterns. Halma's link capacity can be further improved using multi-RF-chain receivers, which can exploit diversity to reduce antenna decoding error. Halma can even be integrated with MIMO spatial multiplexing, by allowing such receivers to simultaneously decode multiple streams of data, sent through different groups of transmit antennas. The throughput/energy tradeoffs in such mechanisms, and their integration with energy-efficient MIMO MAC, will be left for our future exploration. Besides WiFi, we remark that Halma marks a first step in bringing multi-antenna benefits to ZigBee sensors without adding costly RF modules.

8. CONCLUSION

We have explored the feasibility of bringing multi-antenna benefits to single RF-chain wireless devices. Our findings are synthesized in a practical cross-layer design, Halma, that uses antenna index to carry extra bits and adaptive antenna hopping to ensure robustness/efficiency of communication. Halma's modulation/decoding components are simple and built from existing WiFi/ZigBee modules. By integrating antenna hopping with the inherent modulation structures of such practical wireless systems, Halma is able to achieve multiple folds of capacity gain – even higher than existing theoretical prediction [6]. Thus, Halma represents a viable and effective means of realizing multi-antenna networking between energy-constrained wireless devices.

Acknowledgement

The work reported in this paper was supported in part by the NSF under Grant CNS-1318292, CNS-1343363, CNS-1350039 and CNS-1404613.

9. REFERENCES

[1] Cambridge Wireless and ICT KTN, "Positioning Paper: RF Front-End Technology Challenges," 2012.

[2] D. Halperin, B. Greenstein, A. Sheth, and D. Wetherall, "Demystifying 802.11n Power Consumption," in *Proc. of the International Conference on Power Aware Computing and Systems (HotPower)*, 2010.

[3] M. O. Khan, V. Dave, Y.-C. Chen, O. Jensen, L. Qiu, A. Bhartia, and S. Rallapalli, "Model-Driven Energy-Aware Rate Adaptation," in *Proc. of ACM MobiHoc*, 2013.

[4] C.-Y. Li, C. Peng, S. Lu, and X. Wang, "Energy-based Rate Adaptation for 802.11n," in *Prof. of ACM MobiCom*, 2012.

[5] J. Jeganathan, A. Ghrayeb, L. Szczecinski, and A. Ceron, "Space Shift Keying Modulation for MIMO Channels," *IEEE Transactions on Wireless Communications*, vol. 8, no. 7, 2009.

[6] M. Di Renzo, H. Haas, A. Ghrayeb, S. Sugiura, and L. Hanzo, "Spatial Modulation for Generalized MIMO: Challenges, Opportunities, and Implementation," *Proceedings of the IEEE*, vol. 102, no. 1, 2014.

[7] S. Sur, T. Wei, and X. Zhang, "Halma Source Code," 2014. [Online]. Available: http://xyzhang.ece.wisc.edu

[8] Atmel Corp., "REB233SMAD-EK." [Online]. Available: http://www.atmel.com/tools/reb233smad-ek.aspx

[9] X. Zhang and K. G. Shin, "E-MiLi: Energy-Minimizing Idle Listening in Wireless Networks," in *Proc. of ACM MobiCom*, 2011.

[10] K.-Y. Jang, S. Hao, A. Sheth, and R. Govindan, "Snooze: Energy Management in 802.11n WLANs," in *Proc. of ACM CoNEXT*, 2011.

[11] I. Pefkianakis, C.-Y. Li, and S. Lu, "What is Wrong/Right with IEEE 802.11n Spatial Multiplexing Power Save Feature?" in *Proc. of IEEE ICNP*, 2011.

[12] X. Xie, X. Zhang, and K. Sundaresan, "Adaptive Feedback Compression for MIMO Networks," in *ACM MobiCom*, 2013.

[13] A. Cidon, K. Nagaraj, S. Katti, and P. Viswanath, "Flashback: Decoupled Lightweight Wireless Control," in *Proc. of ACM SIGCOMM*, 2012.

[14] D. Halperin, W. Hu, A. Sheth, and D. Wetherall, "Predictable 802.11 Packet Delivery from Wireless Channel Measurements," in *Proc. of ACM SIGCOMM*, 2011.

[15] A. Khattab, J. Camp, C. Hunter, P. Murphy, A. Sabharwal, and E. W. Knightly, "WARP: a Flexible Platform for Clean-Slate Wireless Medium Access Protocol Design," *SIGMOBILE Mob. Comput. Commun. Rev.*, vol. 12, 2008.

[16] T. Schmid, "GNU Radio 802.15.4 En- and Decoding," UCLA NESL TR-UCLA-NESL-200609-06, Tech. Rep., 2006.

[17] Texas Instrument Inc., "CC2420," 2013. [Online]. Available: http://www.ti.com/lit/ds/symlink/cc2420.pdf

[18] A. Amiri Sani, L. Zhong, and A. Sabharwal, "Directional Antenna Diversity for Mobile Devices: Characterizations and Solutions," in *Proc. of ACM MobiCom*, 2010.

[19] S. Farahani, *ZigBee Wireless Networks and Transceivers*. Elsevier Inc., 2008.

[20] Analog Devices Inc., "Choosing the Correct Switch, Multiplexer, or Protection Product for Your Application," 2011.

[21] M. Driusso, F. Babich, M. Kadir, and L. Hanzo, "OFDM Aided Space-Time Shift Keying for Dispersive Downlink Channels," in *Proc. of IEEE VTC*, 2012.

[22] R. Chang, S.-J. Lin, and W.-H. Chung, "Energy Efficient Transmission over Space Shift Keying Modulated MIMO Channels," *IEEE Transactions on Communications*, vol. 60, no. 10, 2012.

[23] R. Mesleh, H. Haas, S. Sinanovic, C. W. Ahn, and S. Yun, "Spatial Modulation," *IEEE Transactions on Vehicular Technology*, vol. 57, no. 4, 2008.

[24] T. Weissman, "Capacity of channels with action-dependent states," *IEEE Transactions on Information Theory*, vol. 56, no. 11, 2010.

[25] A. Younis, W. H. Thompson, M. D. Renzo, C.-X. Wang, M. A. Beach, H. Haas, and P. M. Grant, "Performance of Spatial Modulation using Measured Real-World Channels," *CoRR*, vol. abs/1305.3437, 2013.

[26] E. Magistretti, O. Gurewitz, and E. W. Knightly, "802.11ec: Collision Avoidance Without Control Messages," in *ACM MobiCom*, 2012.

[27] K. Wu, H. Tan, Y. Liu, J. Zhang, Q. Zhang, and L. Ni, "Side Channel: Bits over Interference," in *Proc. of ACM MobiCom*, 2010.

[28] S. Sanayei and A. Nosratinia, "Antenna Selection in MIMO Systems," *IEEE Communications Magazine*, vol. 42, no. 10, 2004.

[29] ——, "Capacity of MIMO Channels With Antenna Selection," *IEEE Transactions on Information Theory*, vol. 53, no. 11, 2007.

[30] C.-M. Cheng, P.-H. Hsiao, H. T. Kung, and D. Vlah, "Transmit Antenna Selection Based on Link-layer Channel Probing," in *Proc. of IEEE WoWMoM*, 2007.

[31] H. Yu, L. Zhong, and A. Sabharwal, "Adaptive RF Chain Management for Energy-efficient Spatial-Multiplexing MIMO Transmission," in *ACM/IEEE ISLPED*, 2009.