Frequency Hopping against a Powerful Adversary

Yuval Emek[1] and Roger Wattenhofer[2]

[1] ETH Zurich, Switzerland. emek@tik.ee.ethz.ch.
[2] Microsoft Research, Redmond, WA and ETH Zurich, Switzerland.
wattenhofer@tik.ee.ethz.ch.

Abstract. Frequency hopping is a central method in wireless communication, offering improved resistance to adversarial interference and interception attempts, and easy non-coordinated control in dynamic environments. In this paper, we introduce a new model that supports a rigorous study of frequency hopping in adversarial settings. We then propose new frequency hopping protocols that allow a sender-receiver pair to essentially use the full communication capacity, despite a powerful adversary that can scan and jam a significant amount of the ongoing transmissions.

1 Introduction

The term *frequency hopping (FH)* in wireless communication refers to a century old method [31–33] of rapidly switching the carrier of a transmitted radio signal among many frequency channels. This method offers various advantages in comparison to traditional fixed frequency transmissions: it is highly resistant to narrow-band interference, it is much more difficult to intercept, and it allows for easy non-coordinated control in dynamic environments. Because of these advantages, FH is omnipresent in modern wireless communication standards such as GSM. Nevertheless, state-of-the-art FH schemes typically use "cryptographic heuristics" whose security is not mathematically established and sometimes turns out to be compromised. For example, using an off-the-shelf device that costs less than 100 EUR [27], the FH scheme that lies at the heart of Bluetooth can be breached within less than a second [17].

In this paper, we hope to bring the state of analysis of FH to the next level. In particular, we ask ourselves what kind of interference a FH protocol can withstand on an information theoretic level (without making any cryptographic assumptions). It turns out that the right tools enable us to design a FH protocol that can cope with adversarial interference, where the adversary can not only jam a constant fraction ϕ of the bandwidth, but also intercept the protocol's transmissions with a small delay. The price to pay for implementing this protocol is a constant additive overhead on the size of each transmitted message. Surprisingly, our protocol manages to utilize the bandwidth up to that ϕ fraction.

1.1 Model

The Cast. Consider the setting of a uni-directional wireless communication from Alice (the transmitter) to Bob (the receiver) in an adversarial environment. There are n available *channels* and in each *round* $t \in \mathbb{Z}_{>0}$, Alice chooses a single channel $a(t) \in [n]$ over which she transmits her message and Bob chooses a single channel $b(t)$ on which he listens; $b(t) = a(t)$ is a necessary condition for Bob to receive Alice's message in round t.

Eve (the adversary) wishes to disturb the communication from Alice to Bob. In each round $t \in \mathbb{Z}_{>0}$, Eve chooses to *jam* a channel subset $E(t) \subset [n]$ of size at most ϕn, $0 < \phi < 1$: if $b(t) \in E(t)$, then Bob does not receive Alice's message even if $b(t) = a(t)$. We distinguish between two types of jamming, differing in the exact effect that $b(t) \in E(t)$ has on Bob's input in round t: *overwriting* means that Bob receives a message that was tailored by Eve which may be confused with Alice's messages; *blocking* means that Bob receives *static noise* which in particular, indicates to Bob that he did not receive Alice's message. Eve is called an *overwriting* (respectively, *blocking*) adversary if her jamming capabilities are suited for overwriting (resp., blocking) Alice's transmissions. For completeness of the model, we assume that if Bob listens on a wrong channel which is not jammed by Eve, i.e., $b(t) \notin E(t) \cup \{a(t)\}$, then he also receives static noise. Note that Alice does not get any feedback regarding the channel on which Bob listened or the actual message he received (if any).

In attempt to avoid Eve's channel jamming, Alice and Bob must use randomness in their channel choices.[3] This should be done in a coordinated fashion to ensure, above everything else, that they both choose the same channel. For that purpose, they both have access to a total of s *shared* random bits generated (once) prior to round 1. Alice can also generate as many *private* random bits as she needs in each round; these cannot be (directly) accessed by Bob, however, Alice may append to each message she transmits up to k additional bits that can be used to communicate some information regarding her (private) random choices. In fact, since the actual content of Alice's messages is irrelevant to the current paper, we shall subsequently consider these k bits as Alice's (whole) message, so in what follows we assume that all messages are of size k.

The setting described so far is trivial to cope with if Eve is an *oblivious* adversary: Alice and Bob can simply follow a random permutation of the channels (assuming that s is sufficiently large to support this random choice, i.e., $s = \Omega(n \log n)$). However, in our model Eve also enjoys the benefit of some sort of *delayed adaptiveness*. It is assumed that Eve can *scan* the spectrum and extract the channel $a(t)$ over which Alice transmitted in round t, but this information is revealed to Eve with a certain *lag* λ, that is, in round $t + \lambda$.

Notice that according to our basic model, Eve's scanning reveals the channel $a(t)$ over which Alice transmitted in round t, but it does not reveal the content of the transmitted message. (Although this issue is abstracted away in our model,

[3] As usual, we assume that Eve knows Alice and Bob's protocol, but not their random bits.

note that it is a valid assumption in settings with encrypted messages.) A variant of the model in which Eve's scanning reveals both the channel and the message content, referred to as *enhanced scanning*, is discussed in Sec. 4.

Grouping the Parameters. We refer to Alice and Bob's channel choosing strategy as an (n, s, k)-*FH protocol*, where n denotes the number of channels, s denotes the number of shared random bits, and k denotes the size (in bits) of the messages. Eve is referred to as a (λ, ϕ)-*adversary*, where λ denotes the delayed adaptiveness lag and ϕ denotes the fraction of channels she jams. It is important to point out that these five parameters may exhibit inter-dependencies (in particular, we shall express s and λ as functions of n), however, unless stated otherwise, they do not grow as a function of the execution length.

Round t is said to be *successful* if $b(t) = a(t) \notin E(t)$, namely, if Bob listens on the channel over which Alice transmits and this channel is not jammed by Eve. The quality of a FH protocol is measured in terms of the fraction of successful rounds captured by the probability that an arbitrary round is successful. Clearly, no FH protocol can guarantee a success probability larger than $1 - \phi$; this is demonstrated already by an (oblivious) adversary that in every round t, chooses $E(t)$ uniformly at random out of all channel subsets of size ϕn.[4] Therefore, at best, we can hope for FH protocols that guarantee success probability close to $1 - \phi$.

Resilience. Formally, an (n, s, k)-FH protocol is said to be ϵ-*resilient* to blocking/overwriting (λ, ϕ)-adversaries if round t is successful with probability at least $1 - \phi - \epsilon$ for every $t \in \mathbb{Z}_{>0}$ against any blocking/overwriting (λ, ϕ)-adversary, respectively. Note that the requirement on the success probability should hold, in particular, as t goes to infinity (fixing all other parameters). This can be thought of as requiring that the guarantees of the FH protocol hold for infinitely long executions, even though all other parameters (including the number s of shared random bits) are finite.

Motivation. The role of the shared random bits is similar to that of a *secret key* in cryptographic systems, generated and exchanged between the collaborating parties before the execution commences. Under our model, the situation is clearly hopeless without shared random bits: if $s = 0$, then Eve knows everything Bob knows already in round 1 and can easily jam the communication. On the other hand, if Alice and Bob have access to infinitely many shared random bits and in particular, can use fresh $\lceil \lg n \rceil$ shared random bits per round,[5] then they can trivially choose $a(t) = b(t) \in [n]$ uniformly at random in every round t, thus ensuring an optimal success probability of $1 - \phi$. To a large extent, the challenge

[4] By Yao's principle, the existence of an oblivious probabilistic adversary that guarantees a success probability of at most $1 - \phi$ against all FH protocols implies that for every FH protocol, there exists an oblivious deterministic adversary with the same guarantee.

[5] We use $\lg x$ to denote $\log_2 x$.

in this paper is to to deal with the case of a finite, yet positive, number s of shared random bits, while trying to keep s small (as a function of n).

One may wonder whether the delayed adaptiveness feature of our model can be justified in practical applications. To that end, note that with dedicated hardware, Eve can scan all n channels, however, extracting the information regarding the channels over which Alice transmitted from the perceived signals is a difficult challenge, likely to incur a significant delay. Moreover, in practical FH scenarios, the spectrum is usually shared between many concurrently communicating Alice-Bob pairs (e.g., secondary users in cognitive radio networks [24]), thus adding another level of complexity to the challenge of obtaining the FH channels used by one specific pair.

1.2 Related Work

Several people are credited with inventing FH. In 1903 Nikola Tesla was granted two U.S. patents [31, 32], where in the second patent, he states: *"To overcome [several drawbacks such as electrical disturbance] and to enable a great number of transmitting and receiving stations to be operated selectively and exclusively and without any danger of the signals or messages being disturbed, intercepted, or interfered with in any way is the object of my present invention."* Jonathan Zenneck [33] claimed in his textbook on wireless telegraphy that the newly founded company Telefunken tested FH around the same time.

The first applications of FH were probably for military purposes. It is reported that the German Reichswehr used FH during World War I to prevent eavesdropping by British forces. During World War II, FH was already pretty common, e.g. in a system called SIGSALY that provided a secure communication infrastructure between Roosevelt and Churchill. Perhaps the most well known FH related martial invention was that of star actress Hedy Lamarr (Markey) and composer George Antheil for preventing the detection of radio guided torpedoes [21]. Nowadays, FH is used by essentially all military radio systems.

FH is well studied in the context of information and coding theory, e.g. [15, 22, 5]. These studies typically aim to provide algebraic hopping sequences with various properties, such as good Hamming correlation or near linear span. However, to the best of our knowledge, this body of work does not deal with adversarial interference.

In contrast, the wireless algorithms community has recently developed an increasing interest in adversarial jamming. Often, the jammer must live on a limited energy budget, which may [16] or may not [13] be known. Dolev et al. [9] studied jamming in the context of multi-channel gossip and presented tight bounds for the ϵ-gossip problem, where the adversary may jam 1 frequency per round. They also study a setting allowing the nodes to exchange authenticated messages despite a malicious adversary that can cause collisions and spoof messages [10], and present new bounds on broadcasting [11]. Another line of work focuses on the *bootstrap problem* where nodes have to find each other despite adversarial jammers [23, 8, 4]. Awerbuch et al. [3] present a MAC protocol for single-hop networks that is provably robust to an adaptive adversary that can

jam (in a blocking style) a $(1-\epsilon)$-fraction of the rounds. This work was later extended to self-stabilization [29, 30]. In [2], the adversary controls both packet injections and jamming, according to a leaky bucket process. Richa et al. [28] recently introduced a reactive jammer that can in addition learn from the protocol history. Hopping sequences with cryptographic guarantees on the resilience to adversarial jamming is studied, e.g., in [19].

Due to their asymptotic approach, theoretical works are typically deemed successful once they manage to exploit a constant fraction of the available communication capacity. In contrast, wireless protocol designers are rarely willing to sacrifice a constant fraction of the precious capacity for protocol overhead. In that regard, we would like to emphasize that our protocols use the available communication capacity up to an ϵ-fraction that can be made arbitrarily small.

1.3 Our Results

Our main technical contribution is a FH protocol that guarantees success probability near $1-\phi$ with constant size messages, logarithmically many shared random bits, and a logarithmic lag. This protocol is suitable for any constant $0 < \phi < 1$ if Eve is a blocking adversary; and for any constant $0 < \phi < 1/16$ if Eve is an overwriting adversary.

We then turn to study the enhanced scanning variant of the model, where the content of Alice's messages is revealed to Eve together with the channel over which these messages were transmitted. In this variant we prove that resilience cannot be achieved as long as the adaptiveness lag is bounded. On the other hand, we show that if the lag *grows* logarithmically with time, then our FH protocol works even when Eve enjoys the benefit of enhanced scanning.

1.4 Techniques

Our FH protocols are inspired by pseudo-random generators à la Impagliazzo and Zuckerman [14]. The sequence of channels over which Alice transmits corresponds to a random walk on an n-vertex constant degree expander. On the one hand, this sequence seems sufficiently random to fool Eve; on the other hand, Bob only needs a constant number of bits per round in order to follow Alice's choices. Since a ϕ-fraction of Alice's messages are doomed to be lost, she encodes her transmissions using a family of error-correcting codes with suitably chosen parameters.

In contrast to the method of Impagliazzo and Zuckerman, where the subset of bad vertices is fixed, we have to deal with an adaptive adversary that dynamically changes the bad vertex subset. This issue is handled in our analysis through a careful examination of the spectral properties of the underlying expander.

2 Preliminaries

In this section, we describe the main ingredients used in the design of our FH protocols, namely, expander graphs and error-correcting codes.

Ramanujan Graphs. Consider some n-vertex d-regular connected non-bipartite graph G. Let $A \in \{0,1\}^{n\times n}$ be G's adjacency matrix and let $W = \frac{1}{d}A$ be the corresponding *walk* matrix. Since W is symmetric, it has n real eigenvalues $\omega_1 \geq \cdots \geq \omega_n$, and since G is d-regular, connected, and non-bipartite, we know that $1 = \omega_1 > \omega_2 \geq \cdots \geq \omega_n > -1$. Moreover, the all 1s vector $\mathbf{1}$ is an eigenvector of W of eigenvalue $\omega_1 = 1$, thus the stationary distribution of the random walk w is uniform in $[n]$.

Let $\omega(G) = \max_{2\leq i\leq n}\{|\omega_i|\} = \max\{\omega_2, |\omega_n|\}$. The parameter $\omega(G)$ captures some important properties of the graph G, and in particular, the speed of convergence of a random walk to the stationary distribution. This is cast in the following lemma which is a well known fact in spectral graph theory (see, e.g., [7]).

Lemma 1. *Let w be a random walk in an n-vertex regular connected non-bipartite graph G and let $\boldsymbol{w}_t$ be the distribution vector of w after t steps. Then for every $i \in [n]$ and $t \in \mathbb{Z}_{>0}$, we have $\left|\boldsymbol{w}_t(i) - \frac{1}{n}\right| \leq \omega(G)^t$. Note that this inequality holds regardless of the initial distribution $\boldsymbol{w}_0$.*

The graphs in an infinite family $\mathcal{G}$ of d-regular connected non-bipartite graphs are called *expanders* if they all have a small $\omega(G)$, that is, if there exists some constant $0 < c < 1$ such that $\omega(G) \leq c$ for all graphs $G \in \mathcal{G}$. In particular, the graphs in $\mathcal{G}$ are said to be *Ramanujan graphs* (a.k.a. *Ramanujan expanders*) if they all satisfy $\omega(G) \leq 2\sqrt{d-1}/d$ [18]. The Alon-Boppana theorem (cf. [26]) essentially states that Ramanujan graphs are the best possible expanders in terms of their small $\omega(G)$.

Theorem 2 ([18, 20, 25]). *For every prime power q and integer $n_0 > 0$, there exist an integer $n = \Theta(n_0)$ and an explicitly constructable n-vertex Ramanujan graph of degree $d = q + 1$.*

Error-Correcting Codes. An *error-correcting code* C over an alphabet Σ is an injective mapping $C : \Sigma^m \to \Sigma^n$, where m and n, $m < n$, are called the *dimension* and the *length* of the code, respectively. We refer to the $|\Sigma|^m$ strings in the image of C as *codewords*. The *minimum distance* of C is the minimum Hamming distance between any two codewords. The ratio of the minimum distance to the length, referred to as the *relative distance* δ of the code, indicates the quality of the code in terms of the number of errors that can be corrected: any number smaller than $\delta n/2$. The ratio of the dimension to the length, referred to as the *rate* $r = m/n$ of the code, indicates the quality of the code in terms of the number of different messages that can be encoded, also known as the *size* of the code: $|\Sigma|^{rn}$.

Fixing some alphabet Σ of size $|\Sigma| = q$, one typically seeks an infinite family $\mathcal{C}_q$ of codes such that both the relative distance and the rate of every code $C \in \mathcal{C}_q$ are bounded below by some constant. Our construction requires an explicit such family in which the relative distance can be made arbitrarily close to 1 by increasing q (and decreasing the rate). Such a family $\mathcal{C}_q$ is designed, e.g., in [1].

Theorem 3 ([1]). *For every real $\xi > 0$, prime power q, and sufficiently large integer n, there exist a real $0 < r = r(\xi) < 1$ and an explicitly constructable error-correcting code over $GF(q)$ of length n, rate at least r, and relative distance at least $1 - \frac{1}{q} - \xi$.*

3 Resilient FH Protocols

Our goal in this section is to establish the following theorem.

Theorem 4. *Consider some constant real $0 < \phi < 1$ (respectively, $0 < \phi < 1/16$). There exist constant integers $k = k(\phi) > 0$ and $n_0 = n_0(\phi) > 0$ such that for every real $\epsilon > 0$ and integer $n \geq n_0$, there exist an integer $\lambda = O(\log(n/\epsilon))$, an integer $s = O(\log(n/\epsilon))$, and an (n, s, k)-FH protocol with ϵ-resilience to blocking (resp., overwriting) (λ, ϕ)-adversaries.*

The basic protocol, presented in Sec. 3.1 and analyzed in Sec. 3.2, is resilient to overwriting (and hence also blocking) adversaries with $0 < \phi < 1/16$. Section 3.3 is dedicated to tuning up our protocol so that it can cope with the whole range of parameter $0 < \phi < 1$ when restricted to blocking adversaries only.

3.1 The Basic Protocol

In a preprocessing stage, Alice and Bob deterministically construct an n-vertex d-regular Ramanujan graph G as promised by Theorem 2, where $d = d(\phi)$ is a constant integer whose value will be determined later on, and identify the vertices of G with the n channels. Note that Theorem 2 does not promise that such a graph exists for every choice of n, however, by taking a graph G of size $n' > n$, and identifying each channel with either $\lfloor n'/n \rfloor$ or $\lceil n'/n \rceil$ vertices, we do not lose more than an (n/n')-term in the guaranteed success probability, and this can be made arbitrarily small. For the sake of simplicity, we shall subsequently assume that the graph G has exactly n vertices. Since the construction of G is deterministic, we are forced to assume that Eve knows G; this will not affect our analysis.

The Phases. Our protocol relies on two parameters: a constant real $\rho = \rho(\phi)$, $0 < \rho < 1$, and an integer $L = O(\log(n/\epsilon))$; the exact values of these two parameters will be determined later on. The rounds of the execution are partitioned into *phases* indexed by the non-negative integers, where phase $j \in \mathbb{Z}_{\geq 0}$ consists of the first

$$\ell(j) = L + \left\lceil 2 \log_{1/\rho}(j+1) \right\rceil$$

rounds not belonging to any phase $j' < j$. Note that this fully determines the phase to which round t belongs for every $t \in \mathbb{Z}_{>0}$.

Alice's channel choices follow a random walk w in G: The channels used in phase 0, namely, the initial vertex $a(1)$ (chosen uniformly at random) and the first $L-1$ steps of w, are dictated by the $s = \lceil \lg n + (L-1) \lg d \rceil = O(\log(n/\epsilon))$

shared random bits. The steps of w in phase $j+1$, $j \in \mathbb{Z}_{\geq 0}$, are dictated by Alice's private random bits and communicated to Bob via the $\ell(j)$ messages sent in phase j. Recall that some of the messages received by Bob in phase j may be transmitted over channels jammed by Eve; to compensate for that, Alice's messages in phase j are encoded by a carefully designed error-correcting code.

Communicating w's Steps. Using the terminology of Theorem 3, we take $\xi = \xi(\phi)$ and $q = q(\phi)$ to be a constant real, $0 < \xi < 1/4$, and a constant (integer) prime power, respectively, whose exact values will be determined later on. Employing Theorem 3, let C_j be the error-correcting code over $GF(q)$ with length $\ell(j)$, relative distance $\delta \geq 1 - \frac{1}{q} - \xi$, and rate $r \geq r(\xi)$, where $0 < r(\xi) < 1$ is the real promised by the theorem.

Let μ denote the $\ell(j+1) \leq \ell(j) + 1 \leq 2\ell(j)$ steps of the random walk w in phase $j+1$. We set the size of Alice's messages to $k = k(\phi) = \lceil \lg q \rceil$; this allows Alice to encode μ using the error-correcting code C_j and to transmit the resulting codeword in phase j, a single letter of the alphabet $GF(q)$ in each round. For that to work, we must make sure that the size of C_j is sufficiently large to encode μ, i.e., that $q^{r\ell(j)} \geq d^{\ell(j+1)}$, which is guaranteed by requiring that the parameter $q = q(\phi)$ satisfies $q \geq d^{2/r(\xi)} \geq d^{2/r}$, and hence $q^{r\ell(j)} \geq d^{2\ell(j)} \geq d^{\ell(j+1)}$. This completes the description of our FH protocol.

3.2 Analysis of the Basic Protocol

For the sake of simplicity, we will prove that our FH protocol is $O(\epsilon)$-resilient (rather than ϵ-resilient). Our analysis relies on the fact that with probability at least $1 - O(\epsilon)$, all phases admit many successful rounds. To formally state this fact (and establish it), we first need some more definitions.

Successful Phases. We say that phase $j \in \mathbb{Z}_{\geq 0}$ is *successful* — an event denoted by A_j — if less than a $(\delta/2)$-fraction of the rounds in the phase are unsuccessful. Note that this implies that if Bob listened on the right channels, then he can successfully decode the codeword transmitted by Alice in phase j. By induction on j, we conclude that the event $A_0 \wedge \cdots \wedge A_{j-1}$ implies that $b(t) = a(t)$ for every round t in phase j. We are now ready to state the two main lemmas of our analysis.

Lemma 5. *Consider some round $t \in \mathbb{Z}_{>0}$ and let $j \in \mathbb{Z}_{\geq 0}$ be the phase to which this round belongs. Conditioned on the event $A_0 \wedge \cdots \wedge A_{j-1}$, round t is successful with probability at least $1 - \phi - \epsilon$.*

Lemma 6. *The event $A_0 \wedge \cdots \wedge A_j$ holds with probability at least $1 - O(\epsilon)$ for every $j \in \mathbb{Z}_{\geq 0}$.*

The remainder of Sec. 3.2 is dedicated to establishing Lemmas 5 and 6, but first, we should convince ourselves that the correctness of our protocol indeed follows from these two lemmas. To that end, consider some round $t \in \mathbb{Z}_{>0}$ in

phase j and let B denote the event that round t is successful. By Lemmas 5 and 6, we have

$$\begin{aligned}\mathbb{P}(B) \geq& \mathbb{P}\left(B \mid A_0 \wedge \cdots \wedge A_{j-1}\right) \cdot \mathbb{P}\left(A_0 \wedge \cdots \wedge A_{j-1}\right) \\ \geq& (1-\phi-\epsilon) \cdot (1-O(\epsilon)) \geq 1-\phi-O(\epsilon)\end{aligned}$$

as required.

Our first step towards establishing Lemmas 5 and 6 is to observe that

$$\begin{aligned}\mathbb{P}\left(A_0 \wedge \cdots \wedge A_m\right) =& \ \mathbb{P}\left(A_m \mid A_0 \wedge \cdots \wedge A_{m-1}\right) \cdot \mathbb{P}\left(A_0 \wedge \cdots \wedge A_{m-1}\right) \\ =& \ \mathbb{P}\left(A_m \mid A_0 \wedge \cdots \wedge A_{m-1}\right) \cdot \mathbb{P}\left(A_{m-1} \mid A_0 \wedge \cdots \wedge A_{m-2}\right) \cdots \\ & \cdots \mathbb{P}\left(A_1 \mid A_0\right) \cdot \mathbb{P}\left(A_0\right) .\end{aligned}$$

Fixing $F_j = \mathbb{P}(\neg A_j \mid A_0 \wedge \cdots \wedge A_{j-1})$, we have

$$\mathbb{P}\left(A_0 \wedge \cdots \wedge A_m\right) = (1-F_0)\cdots(1-F_m) \geq 4^{-F_0} \cdots 4^{-F_m} = 4^{-\sum_{j=0}^{m} F_j} ,$$

where the inequality holds by ensuring that $F_j \leq 1/2$ for every $j \in \mathbb{Z}_{\geq 0}$.

Bounding $\sum F_j$**.** Lemma 6 will be established by showing that $\exp_4(-\sum_{j=0}^{m} F_j) \geq 1 - O(\epsilon)$, or alternatively, that $\exp_4(\sum_{j=0}^{m} F_j) \leq 1 + O(\epsilon) \iff \sum_{j=0}^{m} F_j \leq \log_4(1+O(\epsilon))$. Since $\log_4(1+x) > x/2$ for all $0 < x < 1$, it suffices to show that

$$\sum_{j=0}^{m} F_j \leq O(\epsilon) \tag{1}$$

Take $d = d(\phi)$ to be sufficiently large to ensure that $\frac{2}{\sqrt{d}} \leq \frac{1/4-\sqrt{\phi}}{2}$, which is possible as ϕ is a constant strictly smaller than $1/16$. The following three *auxiliary constants* play a major role in setting the parameters introduced in Sec. 3.1:

$$\alpha = \alpha(\phi) = \frac{1/2 - 2\sqrt{\phi}}{4\sqrt{\phi}+1} , \quad \beta = \beta(\phi) = (1+\alpha)\left(\sqrt{\phi} + \frac{2}{\sqrt{d}}\right) ,$$

$$\gamma = \gamma(\phi) = \frac{1}{\log_4\left(\frac{1}{\beta}\right)} .$$

Since $\sqrt{\phi} < 1/4$, it follows that $0 < \alpha < 1/2$. Moreover, we have $0 < \beta \leq \left(1 + \frac{1/2-2\sqrt{\phi}}{4\sqrt{\phi}+1}\right)\left(\sqrt{\phi} + \frac{1/4-\sqrt{\phi}}{2}\right) = \frac{2\sqrt{\phi}+3/2}{8} < \frac{1}{4}$, and hence $0 < \gamma < 1$.

Recall the parameters $\rho = \rho(\phi)$ and $\xi = \xi(\phi)$ introduced in Sec. 3.1 and fix

$$\rho = 2\beta^{\delta/2} \quad \text{and} \quad \xi = \frac{1-\gamma}{4} .$$

Note that since $0 < \gamma < 1$, it follows that $0 < \xi < 1/4$ as promised. Moreover, by requiring that $q \geq \frac{4}{1-\gamma}$, we ensure that $\delta \geq 1 - \frac{1}{q} - \xi \geq \frac{1+\gamma}{2} > \gamma = \frac{1}{\log_4\left(\frac{1}{\beta}\right)}$. This

implies that $4^{1/\delta} < \frac{1}{\beta}$, hence $0 < \rho = 2\beta^{\delta/2} < 1$ as promised. Fix the integer parameter L introduced in Sec. 3.1 to be

$$L = \max\left\{\left\lceil \log_{1/\rho}\left(\frac{1}{\epsilon}\right)\right\rceil, \left\lceil \log_{\sqrt{d}/2}\left(\frac{\phi n}{\epsilon}\right)\right\rceil, \left\lceil \log_{\sqrt{d}/2}\left(\frac{n}{\alpha}\right)\right\rceil\right\},$$

which yields $L = O(\log(n/\epsilon))$ as promised.

We shall establish (1) by showing that

$$F_j \leq \rho^{\ell(j)}\,; \tag{2}$$

indeed, this suffices since it implies that

$$\sum_{j=0}^{m} F_j \leq \sum_{j=0}^{m} \rho^{L+\lceil 2\log_{1/\rho}(j+1)\rceil} \leq \rho^L \cdot \sum_{j=1}^{\infty} j^{-2} \leq \epsilon \cdot O(1)\,,$$

where the last inequality follows from the requirement that $L \geq \log_{1/\rho}\left(\frac{1}{\epsilon}\right)$.

The Adaptiveness Lag. Recalling that $L = O(\log(n/\epsilon))$, we require that the lag $\lambda = O(\log(n/\epsilon))$ satisfies $\lambda \geq L$. For the sake of the analysis, we think of Eve's scanning as the ability to know in round t, the vertices visited by w in all rounds up to $t - L$. In fact, since the random walk w is memoryless, we may think of Eve as a function that maps the current round index t and the vertex visited by w in round $t - L$ to $E(t)$.

Recall that $\omega(G) \leq \frac{2\sqrt{d-1}}{d} < \frac{2}{\sqrt{d}}$. Since $L \geq \log_{\sqrt{d}/2}\left(\frac{\phi n}{\epsilon}\right)$, we can employ Lemma 1 to conclude that Eve's delayed adaptiveness does not allow her to boost the probability of hitting $a(t)$ by more than an additive term of $\frac{\epsilon}{\phi n}$ per channel, which sums up to an additive term of at most ϵ for all channels jammed by Eve, thus yielding Lemma 5. So, it remains to establish Lemma 6 which is executed by proving that (2) holds. Since $L \geq \log_{\sqrt{d}/2}\left(\frac{n}{\alpha}\right)$ as well, we can employ Lemma 1 once more to establish the following observation.

Observation 7. *Conditioned on Eve's knowledge of the vertex visited by w in round $t - L$, the probability that w visits vertex $i \in [n]$ in round t is at most $\frac{1}{n} + \frac{\alpha}{n} = \frac{1}{n}(1 + \alpha)$.*

Consider some phase $j \in \mathbb{Z}_{\geq 0}$ and let $t_1 \leq \cdots \leq t_\ell$ denote the indices of the $\ell = \ell(j)$ rounds in this phase. Assume that all previous phases were successful, i.e., event $A_0 \wedge \cdots \wedge A_{j-1}$ occurs, so, in particular, Bob knows the random walk w up to the end of phase j, that is, $b(t_h) = a(t_h)$ for every $h \in [\ell]$. Inequality (2) can be established by letting $\mathcal{E} = \{h \in [\ell] \mid E(t_h) \ni a(t_h)\}$ and showing that

$$\mathbb{P}\left(|\mathcal{E}| \geq \delta\ell/2\right) \leq \rho^\ell \tag{3}$$

subject to the assumption that Eve knows $a(t - L)$ at round t.

Given some subset $S \subseteq [\ell]$, let $p_S = \mathbb{P}(\mathcal{E} = S)$. We can express $\mathbb{P}(|\mathcal{E}| \geq \delta\ell/2)$ as

$$\mathbb{P}\left(|\mathcal{E}| \geq \delta\ell/2\right) = \mathbb{P}\left(\bigvee_{S\subseteq[\ell],|S|\geq\delta\ell/2} \mathcal{E} = S\right) \leq \sum_{S\subseteq[\ell],|S|\geq\delta\ell/2} p_S\,.$$

Inequality (3) can now be established by showing that

$$p_S \leq \beta^{\delta\ell/2} \tag{4}$$

for every $S \subseteq [\ell]$, $|S| \geq \delta\ell/2$; indeed, (4) implies that $\mathbb{P}\left(|\mathcal{E}| \geq \delta\ell/2\right) \leq 2^\ell \cdot \beta^{\delta\ell/2} = \left(2\beta^{\delta/2}\right)^\ell = \rho^\ell$ as required.

A Linear Algebraic View. Fix some subset $S \subseteq [\ell]$, $|S| \geq \delta\ell/2$. For every $h \in [\ell]$, let D_h be a diagonal $n \times n$ real matrix defined by setting

$$D_h(i,i) = \begin{cases} 1 & \text{if } h \notin S \\ 1+\alpha & \text{if } h \in S \text{ and } i \in E(t_h) \\ 0 & \text{if } h \in S \text{ and } i \notin E(t_h) \end{cases}$$

for every $i \in [n]$. In other words, D_h is the identity matrix if $h \notin S$; and a matrix having $1+\alpha$ on the diagonal entries corresponding to $E(t_h)$ and 0 elsewhere if $h \in S$. Observe that in the latter case, multiplying a vector by D_h increases all entries corresponding to $E(t_h)$ by a factor of $1+\alpha$ and zeros out all other entries.

Lemma 8. *Denoting the uniform distribution vector on $[n]$ by $\boldsymbol{u} = \frac{1}{n}\mathbf{1}$, we have*

$$p_S \leq \mathbf{1}^T D_\ell W D_{\ell-1} W \cdots D_2 W D_1 \boldsymbol{u}\,.$$

Proof. Let N_h denote the vertex subset $E(t_h)$ if $h \in S$; and the vertex subset $[n] - E(t_h)$ otherwise. Taking B_h to be the event that the random walk w visited (a vertex of) N_h in round t_h, we can express the event $S = \mathcal{E}$ (whose probability we would like to bound) as $B_1 \wedge \cdots \wedge B_\ell$. By Observation 7, the i^{th} entry $D_1\boldsymbol{u}(i)$ of the vector $D_1\boldsymbol{u}$ bounds from above the probability that event B_1 occurred and w visits vertex i in round t_1, given that w is in its stationary distribution $\boldsymbol{u}$ in the beginning of the phase. Employing Observation 7 again, we notice by induction on h that

$$D_h W D_{h-1} \cdots W D_1 \boldsymbol{u}(i)$$

serves as an upper bound on the probability that event $B_1 \wedge \cdots \wedge B_h$ occurred and w visits vertex i in round t_h, given that w is in its stationary distribution $\boldsymbol{u}$ in the beginning of the phase. The assertion follows as multiplying by $\mathbf{1}^T$ simply sums up the entries.

Lemma 8 allows us to complete the proof of Lemma 6 by linear algebraic arguments; indeed, we shall establish (4) by showing that

$$\mathbf{1}^T D_\ell W D_{\ell-1} W \cdots D_2 W D_1 \boldsymbol{u} \leq \beta^{\delta\ell/2}\,. \tag{5}$$

based solely on the definition of the matrices $D_1, \ldots, D_\ell$ and on the assumption that W is the walk matrix of a Ramanujan graph. To that end, observe that

$$\begin{aligned} \mathbf{1}^T D_\ell W D_{\ell-1} W \cdots D_2 W D_1 \boldsymbol{u} &= \mathbf{1}^T D_\ell W D_{\ell-1} W \cdots D_2 W D_1 W \boldsymbol{u} \\ &\leq \|\mathbf{1}\| \cdot \|D_\ell W\| \cdots \|D_1 W\| \cdot \|\boldsymbol{u}\| \\ &= \|D_\ell W\| \cdots \|D_1 W\|\,, \end{aligned}$$

where $\|\boldsymbol{v}\| = \sqrt{\sum_{i=1}^{n} \boldsymbol{v}(i)^2}$ denotes the ℓ_2 norm of vector $\boldsymbol{v}$, $\|M\| = \max_{\boldsymbol{v}\in\mathbb{R}^n-\{\mathbf{0}\}} \frac{\|M\boldsymbol{v}\|}{\|\boldsymbol{v}\|}$ denotes the induced norm of matrix M, and the inequality follows from Cauchy-Schwarz and from some well known properties of the induced matrix norm (see, e.g., [6]).

Bounding $\|D_h W\|$**.** Since W is symmetric, we know that $\|W\| = \max_{i\in[n]} |\omega_i| = 1$, and since D_h is the identity matrix for every $h \notin S$, it follows that

$$\|D_\ell W\| \cdots \|D_1 W\| = \prod_{h\in S} \|D_h W\| \ .$$

Recalling that $|S| \geq \delta\ell/2$, inequality (5), and hence, also Lemma 6, are established due to Lemma 9, whose proof is deferred to the full version.

Lemma 9. *The walk matrix W satisfies $\|D_h W\| \leq \beta$ for every $h \in S$.*

3.3 Extending the Range of Parameter ϕ

Our goal in this section is to adapt the FH protocol presented in Sec. 3.1 and the analysis presented in Sec. 3.2 to blocking adversaries while allowing for any constant $0 < \phi < 1$. The main observation en route to this adaptation is that an error-correcting code that can recover from up to k errors, can alternatively recover from *wiping-off* up to $2k$ letters.

More formally, given some alphabet Σ and a word $u \in \Sigma^n$, let $\mathcal{B}_d(u)$ be the set of all words that can be obtained from u by replacing less than d letters with the designated letter $\flat \notin \Sigma$. In other words, $v \in \mathcal{B}_d(u) \subseteq (\Sigma \cup \{\flat\})^n$ if and only if v disagrees with u on less than d entries in which v has the designated letter $\flat$. The proof of the following observation is deferred to the full version.

Observation 10. *If C is an error-correcting code of length n and minimum distance d, then $\mathcal{B}_d(u) \cap \mathcal{B}_d(v) = \emptyset$ for every two codewords u, v of C.*

The application of Observation 10 is rather straightforward: We can use the error-correcting code C to recover from any number smaller than d of wiped-off letters. In the context of our FH protocol, Alice and Bob can recover from any number smaller than d of blocked rounds. So, except from adjusting some of the parameters, we use here the same protocol that we used against overwriting adversaries, only that this time, Bob can reconstruct the codeword that Alice transmitted in phase j as long as the fraction of unsuccessful rounds is smaller than the relative distance of the code (rather than half the relative distance).

Adjusting the Parameters. The FH protocol presented in Sec. 3.1 and analyzed in Sec. 3.2 relies on the parameters d, ρ, L, ξ, q, δ, and r, and on the three auxiliary constants α, β, and γ. We will use the primed versions d', ρ', L', ξ', q', δ', r' and α', β', γ' to describe the adaptation of this protocol to blocking adversaries. The reader is encouraged to read the remainder of this section in conjunction with Sec. 3.1 and 3.2.

Recall that in the context of overwriting adversaries, we assumed that ϕ is a constant satisfying $0 < \phi < 1/16$ and chose d and $0 < \alpha < 1/2$ so that $0 < \beta = (1+\alpha)\left(\sqrt{\phi} + \frac{2}{\sqrt{d}}\right) < 1/4$. In the context of blocking adversaries, we assume that ϕ is a constant satisfying $0 < \phi < 1$ and choose the parameter $d' = d'(\phi)$ and the auxiliary constant $0 < \alpha' = \alpha'(\phi) < 1$ so that

$$0 < \beta' = \beta'(\phi) = (1+\alpha')\left(\sqrt{\phi} + \frac{2}{\sqrt{d'}}\right) < 1\,.$$

Let $H(x) = -x\lg(x) - (1-x)\lg(1-x) = H(1-x)$ be the binary entropy function defined for every $0 < x < 1$. Observe that $\lim_{x\to 1-} 2^{H(x)}\beta'^{x} = \beta'$ and take $\gamma' = \gamma'(\phi)$ to be the smallest real $1/2 \leq \gamma' < 1$ such that $2^{H(\gamma')}\beta'^{\gamma'} \leq \frac{\beta'+1}{2}$. Let $\xi' = \xi'(\phi) = \frac{1-\gamma'}{2}$ and let $q' = q'(\phi)$ be the smallest prime power that satisfies

$$q' \geq \max\left\{\frac{2}{1-\gamma'}, d'^{2/r(\xi')}\right\},$$

where $0 < r(\xi') < 1$ is the real promised by Theorem 3. The error-correcting codes C'_j over $GF(q')$ we use have rate $r' \geq r(\xi')$ and relative distance $\delta' \geq 1 - \frac{1}{q'} - \xi'$. Finally, let $\rho' = \rho'(\phi) = 2^{H(\gamma')}\beta'^{\gamma'}$ and

$$L' = \max\left\{\left\lceil \log_{1/\rho'}\left(\frac{1}{\epsilon}\right)\right\rceil, \left\lceil \log_{\sqrt{d'}/2}\left(\frac{\phi n}{\epsilon}\right)\right\rceil, \left\lceil \log_{\sqrt{d'}/2}\left(\frac{n}{\alpha'}\right)\right\rceil\right\}.$$

Modified Analysis. Using the adapted parameters, the analysis presented in Sec. 3.2 carries over quite smoothly. The one part that does require some changes is that involving inequalities (3) and (4) and the transition between them. Recalling that a phase is now considered to be successful if less than a δ'-fraction of its rounds are unsuccessful, we rewrite (3) as

$$\mathbb{P}\left(|\mathcal{E}| \geq \delta'\ell\right) \leq \rho'^{\ell}\,, \tag{6}$$

(again, subject to the assumption that Eve knows $a(t-L')$ at round t). So, our goal is to prove that (6) follows from $p_S \leq \beta'^{\delta'\ell}$ (the equivalent of (4)) for every $S \subseteq [\ell]$, $|S| \geq \delta'\ell$.

As in Sec. 3.2, we express $\mathbb{P}(|\mathcal{E}| \geq \delta'\ell)$ as

$$\mathbb{P}\left(|\mathcal{E}| \geq \delta'\ell\right) = \mathbb{P}\left(\bigvee_{S\subseteq[\ell],|S|\geq\delta'\ell} \mathcal{E} = S\right) \leq \sum_{S\subseteq[\ell],|S|\geq\delta'\ell} p_S\,,$$

only that this time, we use the fact that $\sum_{k=\lceil xn\rceil}^{n}\binom{n}{k} \leq 2^{H(x)\cdot n}$ for every $n \geq 1$ and $1/2 \leq x < 1$ (see, e.g., [12]) to bound the number of subsets S that should be accounted for. Specifically, we get

$$\mathbb{P}\left(|\mathcal{E}| \geq \delta'\ell\right) \leq 2^{H(\delta')\cdot\ell}\beta'^{\delta'\ell} = \left(2^{H(\delta')}\beta'^{\delta'}\right)^{\ell}.$$

This concludes our proof as $\delta' \geq 1 - \frac{1}{q'} - \xi' \geq \gamma'$, and hence $2^{H(\delta')}\beta'^{\delta'} \leq 2^{H(\gamma')}\beta'^{\gamma'} = \rho'$.

4 Enhanced Scanning

The FH protocols developed in Sec. 3 are ϵ-resilient to (blocking and overwriting) (λ, ϕ)-adversaries that have access in round $t + \lambda$ to the channel $a(t)$ over which Alice transmitted in round t, but not to the actual content $m(t)$ of Alice's message. We now turn our attention to adversaries with *enhanced scanning*, namely, both $a(t)$ and $m(t)$ are revealed to Eve in round $t + \lambda$. On the negative side, we prove that no FH protocol can be resilient to such adversaries as long as we stick to the model introduced in Sec. 1.1, requiring that the lag λ is fixed with respect to the time t. On the positive side, we show that a FH protocol with resilience to enhanced scanning adversaries does exist if the lag λ grows logarithmically with t. Due to space limitations, these proofs are deferred to the full version.

References

1. Alon, N., Bruck, J., Naor, J., Naor, M., Roth, R.M.: Construction of asymptotically good low-rate error-correcting codes through pseudo-random graphs. IEEE Transactions on Information Theory 38, 509–516 (1992)
2. Anantharamu, L., Chlebus, B.S., Kowalski, D.R., Rokicki, M.A.: Medium access control for adversarial channels with jamming. In: Proceedings of the 18th international conference on Structural information and communication complexity. pp. 89–100. SIROCCO, Berlin, Heidelberg (2011)
3. Awerbuch, B., Richa, A., Scheideler, C.: A Jamming-Resistant MAC Protocol for Single-Hop Wireless Networks. In: Proc. 27th Symposium on Principles of Distributed Computing (PODC) (2008)
4. Azar, Y., Gurel-Gurevich, O., Lubetzky, E., Moscibroda, T.: Optimal discovery strategies in white space networks. In: 19th Annual European Symposium of Algorithms, Saarbrücken, Germany. pp. 713–722 (2011)
5. Chu, W., Colbourn, C.: Optimal frequency-hopping sequences via cyclotomy. IEEE Transactions on Information Theory 51(3), 1139 –1141 (March 2005)
6. Demmel, J.W.: Applied numerical linear algebra. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA (1997)
7. Diaconis, P., Stroock, D.: Geometric Bounds for Eigenvalues of Markov Chains. The Annals of Applied Probability 1(1), 36–61 (1991)
8. Dolev, S., Gilbert, S., Guerraoui, R., Kuhn, F., Newport, C.: The wireless synchronization problem. In: Proceedings of the 28th ACM symposium on Principles of distributed computing. pp. 190–199. PODC, New York, NY, USA (2009)
9. Dolev, S., Gilbert, S., Guerraoui, R., Newport, C.: Gossiping in a Multi-Channel Radio Network (An Oblivious Approach to Coping With Malicious Interference). In: Proc. 21st Ann. Conference on Distributed Computing (DISC) (2007)
10. Dolev, S., Gilbert, S., Guerraoui, R., Newport, C.: Secure Communication over Radio Channels. In: Proc. 27th ACM Symposium on Principles of Distributed Computing (PODC). pp. 105–114 (2008)
11. Dolev, S., Gilbert, S., Khabbazian, M., Newport, C.: Leveraging channel diversity to gain efficiency and robustness for wireless broadcast. In: Proceedings of the 25th international conference on Distributed computing. pp. 252–267. DISC'11, Berlin, Heidelberg (2011)

12. Flum, J., Grohe, M.: Parameterized Complexity Theory, vol. 7. Springer (2006)
13. Gilbert, S., Guerraoui, R., Newport, C.: Of Malicious Motes and Suspicious Sensors. In: Proc. 10th Conference on Principles of Distributed Systems (OPODIS) (2006)
14. Impagliazzo, R., Zuckerman, D.: How to recycle random bits. In: Proceedings of the 30th Annual Symposium on Foundations of Computer Science. pp. 248–253. Washington, DC, USA (1989)
15. Knopp, R., Humblet, P.: On coding for block fading channels. IEEE Transactions on Information Theory 46(1), 189 –205 (Jan 2000)
16. Koo, C.Y., Bhandari, V., Katz, J., Vaidya, N.H.: Reliable Broadcast in Radio Networks: the Bounded Collision Case. In: Proc. 25th ACM Symposium on Principles of Distributed Computing (PODC) (2006)
17. Köppel, S.: Bluetooth jamming. Bachelor's Thesis supervised by Michael König and Roger Wattenhofer, ETH Zurich (2013)
18. Lubotzky, A., Phillips, R., Sarnak, P.: Ramanujan graphs. Combinatorica 8, 261–277 (1988)
19. Mansour, I., Chalhoub, G., Quilliot, A.: Security architecture for wireless sensor networks using frequency hopping and public key management. In: ICNSC. pp. 526–531. IEEE (2011)
20. Margulis, G.A.: Explicit group-theoretic constructions of combinatorial schemes and their applications in the construction of expanders and concentrators. Problemy Peredachi Informatsii 24(1), 51–60 (1988)
21. Markey, H.K., Antheil, G.: Secret communication system (1942), U.S. Patent 2292387
22. Medard, M., Gallager, R.: Bandwidth scaling for fading multipath channels. IEEE Transactions on Information Theory 48(4), 840 –852 (April 2002)
23. Meier, D., Pignolet, Y.A., Schmid, S., Wattenhofer, R.: Speed Dating despite Jammers. In: 5th IEEE International Conference on Distributed Computing in Sensor Systems (DCOSS), Marina del Rey, California, USA (June 2009)
24. Mitola, J., Maguire, G.Q.: Cognitive radio: making software radios more personal. Personal Communications, IEEE 6(4), 13–18 (Aug 1999)
25. Morgenstern, M.: Existence and Explicit Constructions of q + 1 Regular Ramanujan Graphs for Every Prime Power q. Journal of Combinatorial Theory, Series B 62(1), 44–62 (1994)
26. Nilli, A.: On the second eigenvalue of a graph. Discrete Math. 91, 207–210 (August 1991)
27. Project Ubertooth. http://ubertooth.sourceforge.net/
28. Richa, A., Scheideler, C., Schmid, S., Zhang, J.: Competitive and fair medium access despite reactive jamming. In: Distributed Computing Systems (ICDCS), 2011 31st International Conference on. pp. 507 –516 (June 2011)
29. Richa, A., Scheideler, C., Schmid, S., Zhang, J.: A jamming-resistant mac protocol for multi-hop wireless networks. In: Proceedings of the 24th international conference on Distributed computing. pp. 179–193. DISC (2010)
30. Richa, A., Scheideler, C., Schmid, S., Zhang, J.: Self-stabilizing leader election for single-hop wireless networks despite jamming. In: Proceedings of the Twelfth ACM International Symposium on Mobile Ad Hoc Networking and Computing. MobiHoc, New York, NY, USA (2011)
31. Tesla, N.: Method of signaling (1903), U.S. Patent 723188
32. Tesla, N.: System of signaling (1903), U.S. Patent 725605
33. Zenneck, J.: Leitfaden der drahtlosen Telegraphie. Enke, Stuttgart, Germany (1909)