

Optimal whitespace synchronization strategies

Yossi Azar* Ori Gurel-Gurevich† Eyal Lubetzky‡ Thomas Moscibroda§

Abstract

The whitespace-discovery problem describes two parties, Alice and Bob, trying to establish a communication channel over one of a given large segment of whitespace channels. Subsets of the channels are occupied in each of the local environments surrounding Alice and Bob, as well as in the global environment between them (Eve). In the absence of a common clock for the two parties, the goal is to devise time-invariant (stationary) strategies minimizing the synchronization time. This emerged from recent applications in discovery of wireless devices.

We model the problem as follows. There are N channels, each of which is open (unoccupied) with probability p_1, p_2, q independently for Alice, Bob and Eve respectively. Further assume that $N \gg 1/(p_1 p_2 q)$ to allow for sufficiently many open channels. Both Alice and Bob can detect which channels are locally open and every time-slot each of them chooses one such channel for an attempted sync. One aims for strategies that, with high probability over the environments, guarantee a shortest possible expected sync time depending only on the p_i 's and q .

Here we provide a stationary strategy for Alice and Bob with a guaranteed expected sync time of $O(1/(p_1 p_2 q^2))$ given that each party also has knowledge of p_1, p_2, q . When the parties are oblivious of these probabilities, analogous strategies incur a cost of a poly-log factor, i.e. $\tilde{O}(1/(p_1 p_2 q^2))$. Furthermore, this performance guarantee is essentially optimal as we show that any stationary strategies of Alice and Bob have an expected sync time of at least $\Omega(1/(p_1 p_2 q^2))$.

1 Introduction

Consider two parties, Alice and Bob, who wish to establish a communication channel in one out of a segment of N possible channels. Subsets of these channels may already be occupied in the local environments of either Alice or Bob, as well as in the global environment in between them whose users are denoted by Eve. Furthermore, the two parties do not share a common clock and hence one does not know for how long (if at all) the other party has already been trying to communicate. Motivated by applications in discovery of wireless devices, the goal is thus to devise time-invariant strategies that ensure fast synchronization with high probability (w.h.p.) over the environments.

We formalize the above problem as follows. Transmissions between Alice and Bob go over three environments: local ones around Alice and Bob and an additional global one in between them, Eve. Let A_i, B_i, E_i for $i = 1, \dots, N$ be the indicators for whether a given channel is open (unoccupied)

*Department of Computer Science, Tel-Aviv University, Tel-Aviv 69978, Israel. Email address: azar@tau.ac.il.

†Microsoft Research, One Microsoft Way, Redmond, WA 98052, USA. Email: origurel@microsoft.com.

‡Microsoft Research, One Microsoft Way, Redmond, WA 98052, USA. Email: eyal@microsoft.com.

§Microsoft Research, One Microsoft Way, Redmond, WA 98052, USA. Email: moscitho@microsoft.com.

in the respective environment. Using local diagnostics Alice knows A yet does not know B, E and analogously Bob knows B but is oblivious of A, E . In each time-slot, each party selects a channel to attempt communication on (the environments do not change between time slots). The parties are said to *synchronize* once they select the same channel i that happens to be open in all environments (i.e., $A_i = B_i = E_i = 1$). The objective of Alice and Bob is to devise strategies that would minimize their expected synchronization time.

For a concrete setup, let A_i, B_i, E_i be independent Bernoulli variables with probabilities p_1, p_2, q respectively for all i , different channels being independent of each other. (In some applications the two parties have knowledge of the environment densities p_1, p_2, q while in others these are unknown.) Alice and Bob then seek strategies whose expected sync time over the environments is minimal.

Example. Suppose that $p_1 = p_2 = 1$ (local environments are fully open) and Alice and Bob use the naive strategy of selecting a channel uniformly over $[N]$ and independently every round. If there are $Q \approx qN$ open channels in the global environment Eve then the probability of syncing in a given round is $Q/N^2 \approx q/N$, implying an expected sync time of about N/q to the very least.

In the above framework it could occur that *all* channels are closed, in which case the parties can never sync; as a result, unless this event is excluded the expected sync time is always infinite. However, since this event has probability at most $(1 - p_1 p_2 q)^N \leq \exp(-N p_1 p_2 q)$ it poses no real problem for applications (described in further details later) where $N \gg 1/(p_1 p_2 q)$. In fact, we aim for performance guarantees that depend only on p_1, p_2, q rather than on N , hence a natural way to resolve this issue is to extend the set of channels to be infinite, i.e. define A_i, B_i, E_i for every $i \in \mathbb{N}$. (Our results can easily be translated to the finite setting with the appropriate exponential error probabilities.)

A *strategy* is a sequence of probability measures $\{\mu^t\}$ over \mathbb{N} , corresponding to a randomized choice of channel for each time-slot $t \geq 1$. Suppose that Alice begins the discovery via the strategy μ_a whereas Bob begins the synchronization attempt at time s via the strategy μ_b . Let X_t be the indicator for a successful sync at time t and let X be the first time Alice and Bob sync, that is

$$\mathbb{P}(X_t = 1 \mid A, B, E) = \sum_j \mu_a^t(j) \mu_b^{s+t}(j) A_j B_j E_j, \quad (1.1)$$

$$X = \min\{t : X_t = 1\}. \quad (1.2)$$

The choice of μ_a, μ_b aims to minimize $\mathbb{E}X$ where the expectation is over A, B, E as well as the randomness of Alice and Bob in applying the strategies μ_a, μ_b .

Example (fixed strategies). Suppose that both Alice and Bob apply the same pair of strategies independently for all rounds, μ_a and μ_b respectively. In this special case, given the environments A, B, E the random variable X is geometric with success probability $\sum_j \mu_a(j) \mu_b(j) A_j B_j E_j$, thus the mappings $A \mapsto \mu_a$ and $B \mapsto \mu_b$ should minimize $\mathbb{E}X = \mathbb{E} \left[\left(\sum_j \mu_a(j) \mu_b(j) A_j B_j E_j \right)^{-1} \right]$.

A crucial fact in our setup is that Alice and Bob have no common clock and no means of telling whether or not their peer is already attempting to communicate (until they eventually synchronize). As such, they are forced to apply a *stationary* strategy, where the law at each time-slot is identical (i.e. $\mu^t \sim \mu^1$ for all t). For instance, Alice may choose a single μ_a and apply it independently in each

step (cf. above example). Alternatively, strategies of different time-slots can be highly dependent, e.g. Bob may apply a periodic policy comprising μ_b^1, \dots, μ_b^n and a uniform initial state $s \in [n]$.

The following argument demonstrates that stationary strategies are essentially optimal when there is no common clock between the parties. Suppose that Alice has some finite (arbitrarily long) sequence of strategies $\{\mu_a^t\}_1^{M_a}$ and similarly Bob has a sequence of strategies $\{\mu_b^t\}_1^{M_b}$. With no feedback until any actual synchronization we may assume that the strategies are non-adaptive, i.e. the sequences are determined in advance. Without loss of generality Alice is joining the transmission after Bob has already attempted some β rounds of communication, in which case the expected synchronization time is $\mathbb{E}_{0,\beta}X$, where $\mathbb{E}_{\alpha,\beta}X$ denotes the expectation of X as defined in (1.1),(1.2) using the strategies $\{\mu_a^{t+\alpha}\}, \{\mu_b^{t+\beta}\}$. Having no common clock implies that in the worst case scenario (over the state of Bob) the expected time to sync is $\max_{\beta} \mathbb{E}_{0,\beta}X$ and it now follows that Bob is better off modifying his strategy into a stationary one by selecting $\beta \in [M_b]$ uniformly at random, leading to an expected synchronization time of $M_b^{-1} \sum_{\beta} \mathbb{E}_{0,\beta}X$.

1.1 Optimal synchronization strategies

Our main result is a recipe for Alice and Bob to devise stationary strategies guaranteeing an optimal expected synchronization time up to an absolute constant factor, assuming they know the environment densities p_1, p_2, q (otherwise the expected sync time is optimal up to a poly-log factor).

Theorem 1. *Consider the synchronization problem with probabilities p_1, p_2, q for the environments A, B, E respectively, and let X denote the expected sync time. The following then holds:*

(i) *There are fixed strategies for Alice and Bob guaranteeing $\mathbb{E}X = O(1/(p_1 p_2 q^2))$, namely:*

- *Alice takes $\mu_a \sim \text{Geom}(p_2 q/6)$ over her open channels $\{i : A_i = 1\}$,*
- *Bob takes $\mu_b \sim \text{Geom}(p_1 q/6)$ over his open channels $\{i : B_i = 1\}$.*

Furthermore, for any fixed $\varepsilon > 0$ there are fixed strategies for Alice and Bob that do not require knowledge of p_1, p_2, q and guarantee $\mathbb{E}X = O\left(\frac{1}{p_1 p_2 q^2} \log^{2+\varepsilon}\left(\frac{1}{p_1 p_2 q}\right)\right) = \tilde{O}\left(\frac{1}{p_1 p_2 q^2}\right)$, obtained by taking $\mu_a(j\text{-th open } A \text{ channel}) = \mu_b(j\text{-th open } B \text{ channel}) \propto 1/(j \log^{1+\varepsilon/2} j)$.

(ii) *The above strategies are essentially optimal as every possible choice of stationary strategies by Alice and Bob satisfies $\mathbb{E}X = \Omega(1/(p_1 p_2 q^2))$.*

Remark. The factor 1/6 in the parameters of the geometric distributions can be fine-tuned to any smaller (or even slightly larger) fixed $\alpha > 0$ affecting the overall expected sync time $\mathbb{E}X$ by a multiplicative constant. See Fig. 1 for a numerical evaluation of $\mathbb{E}X$ for various values of α .

Recall that Alice and Bob must apply stationary strategies in the absence of any common clock or external synchronization device shared by them, a restriction which is essential in many of the applications of wireless discovery protocols. However, whenever a common external clock does happen to be available there may be strategies that achieve improved performance. The next theorem establishes the optimal strategies in this simpler scenario.

Theorem 2. *Consider the synchronization problem with probabilities p_1, p_2, q for the environments A, B, E respectively, and let X denote the expected sync time. If Alice and Bob have access to a common clock then there are non-stationary strategies for them achieving $\mathbb{E}X = O(1/(\min\{p_1, p_2\}q))$. Moreover, this is tight as the expected sync time is always $\Omega(1/(\min\{p_1, p_2\}q))$.*




Figure 1: Synchronization time $\mathbb{E}X$ as in (1.2) normalized by a factor of $p_1 p_2 q^2$ for a protocol using geometric distributions with parameters $\alpha p_i q$ for various values of $0 < \alpha < 1$. Markers represent the average of the expected synchronization time $\mathbb{E}X$ over 10^5 random environments with $n = 10^4$ channels; surrounding envelopes represent a window of one standard deviation around the mean.

1.2 Applications in wireless networking and related work

The motivating application for this work comes from recent developments in wireless networking. In late 2008, the FCC issued a historic ruling permitting the use of unlicensed devices in the unused portions of the UHF spectrum (mainly the part between 512Mhz and 698Mhz), popularly referred to as “whitespaces”. Due to the potential for substantial bandwidth and long transmission ranges, whitespace networks (which are, in a more general context, also frequently referred to as cognitive wireless networks) represent a tremendous opportunity for mobile and wireless communication. One critical equipment imposed by the FCC in its ruling is that whitespace wireless devices must not interfere with incumbents, i.e., the current users of this spectrum (specifically, in the UHF bands, these are TV broadcasters as well as wireless microphones). Hence, these incumbents are considered “primary users” of the spectrum, while whitespace devices are secondary users and are allowed to use the spectrum only opportunistically, whenever no primary user is using it (The FCC mandates whitespace devices to detect the presence of primary users using a combination of sensing techniques and a geo-location database). At any given time, each whitespace device thus has a spectrum map on which some parts are blocked off while others are free to use.

The problem studied in this paper captures exactly the situation in whitespace networks when two nodes A and B seek to discover one another to establish a connection. Each node knows its own free channels on which it can transmit, but it does not know which of these channels may be available at the other node, too. Furthermore, given the larger transmission range in whitespace networks (up to a mile at Wi-Fi transmission power levels), it is likely that the spectrum maps at A and B are similar yet different. For example, a TV broadcast tower is likely to block off a

channel for both A and B , but a wireless microphone — due to its small transmission power — will prevent only one of the nodes from using a channel.

Thus far, the problem of synchronizing/discovery of whitespace nodes has only been addressed when one of the nodes is a fixed access point (AP) and the other node is a client. Namely, in the framework studied in [5] the AP broadcasts on a fixed channel and the client node wishes to scan its local environment and locate this channel efficiently. That setting thus calls for technological solutions (e.g. based on scanning wider channel widths) to allow the client to find the AP channel faster than the approach of searching all possible channels one by one.

To the best of our knowledge, the results in this paper are the first to provide an efficient synchronization scheme in the setting where both nodes are remote clients that may broadcast on any given channel in the whitespace region.

1.3 Related work on Rendezvous games

From a mathematical standpoint, the synchronization problems considered in this paper seem to belong to the field of Rendezvous Search Games. The most familiar problem of this type is known as The Telephone Problem or The Telephone Coordination Game. In the telephone problem each of two players is placed in a distinct room with n telephone lines connecting the rooms. The lines are not labeled and so the players, who wish to communicate with each other, cannot simply use the first line (note that, in comparison, in our setting the channels are labeled and the difficulty in synchronizing is due to the local and global noise).

The optimal strategy in this case, achieving an expectation of $n/2$, is for the first player to pick a random line and continue using it, whereas the second player picks a uniformly random permutation on the lines and try them one by one. However, this strategy requires the players to determine which is the first and which is the second. It is very plausible that such coordination is not possible, in which case we require both players to employ the same strategy.

The obvious solution is for each of them to pick a random line at each turn, which gives an expectation of n turns. It turns out, however, that there are better solutions: Anderson and Weber [4] give a solution yielding an expectation of $\approx 0.8288497n$ and conjecture it's optimality.

To our knowledge, the two most prominent aspects of our setting, the presence of asymmetric information and the stationarity requirement (stemming from unknown start times) have not been considered in the literature. For example, the Anderson-Weber strategy for the telephone problem is not stationary — it has a period of $n - 1$. It would be interesting to see what can be said about the optimal stationary strategies for this and other rendezvous problems. The interested reader is referred to [2,3] and the references therein for more information on rendezvous search games.

2 Analysis of synchronization strategies

2.1 Proof of Theorem 1, upper bound on the sync time

Let μ_a be geometric with mean $(\alpha p_2 q)^{-1}$ over the open channels for Alice $\{i : A_i = 1\}$ and analogously let μ_b be geometric with mean $(\alpha p_1 q)^{-1}$ over the open channels for Bob $\{i : B_i = 1\}$,

where $0 < \alpha < 1$ will be determined later.

Let $J = \min\{j : A_j = B_j = E_j = 1\}$ be the minimal channel open in all three environments. Further let J_a, J_b denote the number of locally open channels prior to channel J for Alice and Bob resp., that is

$$J_a = \#\{j < J : A_j = 1\}, \quad J_b = \#\{j < J : B_j = 1\}.$$

Finally, for some integer $k \geq 0$ let M_k denote the event

$$k \leq \max\{J_a p_2 q, J_b p_1 q\} < k + 1. \quad (2.1)$$

Notice that, by definition, Alice gives probability $(1 - \alpha p_2 q)^{j-1} \alpha p_2 q$ to her j -th open channel while Bob gives probability $(1 - \alpha p_1 q)^{j-1} \alpha p_1 q$ to his j -th open channel. Therefore, on the event M_k we have that in any specific round, channel J is chosen by both players with probability at least

$$(1 - \alpha p_1 q)^{\frac{k+1}{p_1 q}} (1 - \alpha p_2 q)^{\frac{k+1}{p_2 q}} \alpha^2 p_1 p_2 q^2 \geq e^{-4\alpha(k+1)} \alpha^2 p_1 p_2 q^2,$$

where in the last inequality we used the fact that $(1 - x) \geq \exp(-2x)$ for all $0 \leq x \leq \frac{1}{2}$, which will be justified by later choosing $\alpha < \frac{1}{2}$. Therefore, if X denotes the expected number of rounds required for synchronization, then

$$\mathbb{E}[X \mid M_k] \leq e^{4\alpha(k+1)} (\alpha^2 p_1 p_2 q^2)^{-1}. \quad (2.2)$$

On the other hand, J_a is precisely a geometric variable with the rule $\mathbb{P}(J_a = j) = (1 - p_2 q)^j p_2 q$ and similarly $\mathbb{P}(J_b = j) = (1 - p_1 q)^j p_1 q$. Hence,

$$\mathbb{P}(M_k) \leq (1 - p_2 q)^{k/(p_2 q)} + (1 - p_1 q)^{k/(p_1 q)} \leq 2e^{-k}.$$

Combining this with (2.2) we deduce that

$$\mathbb{E}X \leq 2 \sum_k e^{-k} \mathbb{E}[X \mid M_k] \leq 2e^{4\alpha} (\alpha^2 p_1 p_2 q^2)^{-1} \sum_k e^{(4\alpha-1)k} \leq \frac{2e}{\alpha^2 (e^{1-4\alpha} - 1)} (p_1 p_2 q^2)^{-1} \quad (2.3)$$

where the last inequality holds for any fixed $\alpha < \frac{1}{4}$. In particular, a choice of $\alpha = \frac{1}{6}$ implies that $\mathbb{E}X \leq 500 / (p_1 p_2 q^2)$, as required. \blacksquare

Remark 2.1. In the special case where $p_1 = p_2$ (denoting this probability simply by p) one can optimize the choice of constants in the proof above to obtain an upper bound of $\mathbb{E}X \leq 27 / (pq)^2$.

2.2 Strategies oblivious of the environment densities

Observe first that if Alice and Bob multiply the parameters of their geometric distributions (as specified in Part (i) of Theorem 1) by some absolute constant $0 < c < 1$ then the upper bound (2.3) on $\mathbb{E}X$ will increase by a factor of at most some absolute $C > 0$.

With this in mind, fix $\varepsilon > 0$ and consider the following strategies. Every round, Alice guesses $p_2 q$ to be $\exp(-i)$ for $i = 1, 2, \dots$ with probability proportional to $i^{-(1+\varepsilon/2)}$, while Bob does the

same for p_1q . This way, both Alice and Bob successfully guess these parameters (within a factor of e) with probability at least

$$c \log^{-(1+\varepsilon/2)} \left(\frac{1}{p_1q} \right) \log^{-(1+\varepsilon/2)} \left(\frac{1}{p_2q} \right) \geq c \log^{-(2+\varepsilon)} \left(\frac{1}{p_1p_2q} \right)$$

where $c > 0$ is some absolute constant. Hence, restricting our analysis to these rounds only we obtain an expected sync time of at most $O(1/(p_1p_2q^2))$ such rounds and overall

$$\mathbb{E}X = O\left(\frac{1}{p_1p_2q^2} \log^{2+\varepsilon} \left(\frac{1}{p_1p_2q} \right)\right) = \tilde{O}\left(\frac{1}{p_1p_2q^2}\right).$$

An elementary calculation shows that the above strategy is equivalent to having both Alice and Bob choose their corresponding j -th open channel ($j = 1, 2, \dots$) with probability proportional to $j \log^{1+\varepsilon/2} j$. Indeed, one may repeat the arguments from the previous section directly for these strategies and obtain that (in the original notation) for any given round

$$\mathbb{P}(\text{Alice and Bob select } J \mid M_k) \geq cp_1p_2q^2 \left[(k+1)^2 \log^{1+\varepsilon/2} \left(\frac{k+1}{p_1q} \right) \log^{1+\varepsilon/2} \left(\frac{k+1}{p_2q} \right) \right]^{-1}$$

for some absolute constant $c > 0$. From this we then infer that

$$\begin{aligned} \mathbb{E}X &\leq O(1/(p_1p_2q^2)) \sum_k e^{-k} (k+1)^2 \log^{1+\varepsilon/2} \left(\frac{k+1}{p_1q} \right) \log^{1+\varepsilon/2} \left(\frac{k+1}{p_2q} \right) \\ &= O\left(\frac{1}{p_1p_2q^2} \log^{2+\varepsilon} \left(\frac{1}{p_1p_2q} \right)\right), \end{aligned}$$

as argued above. ■

2.3 Proof of Theorem 1, lower bound on the sync time

Theorem 2.2. *Let μ_a, μ_b be the stationary distribution of the strategies of Alice and Bob resp., and let $R = \sum_j \mu_a(j)\mu_b(j)A_jB_jE_j$ be the probability of successfully syncing in any specific round. Then there exists some absolute constant $C > 0$ such that $\mathbb{P}(R < Cp_0p_1q^2) \geq \frac{1}{2}$.*

Proof. Given the environments A, B define

$$S_k^a = \{j : 2^{-k} < \mu_a(j) \leq 2^{-k+1}\}, \quad S_k^b = \{j : 2^{-k} < \mu_b(j) \leq 2^{-k+1}\}.$$

Notice that the variables S_k^a are a function of the strategy of Alice which in turn depends on her local environment A (an analogous statement holds for S_k^b and B). Further note that clearly

$$|S_k^a| < 2^k \quad \text{and} \quad |S_k^b| < 2^k \quad \text{for any } k.$$

Let T_k^a denote all the channels where the environments excluding Alice's (i.e., both of the other environments B, E) are open, and similarly let T_k^b denote the analogous quantity for Bob:

$$T_k^a = \{j \in S_k^a : B_j = E_j = 1\}, \quad T_k^b = \{j \in S_k^b : A_j = E_j = 1\}.$$

Obviously, $\mathbb{E}|T_k^a| < 2^k p_2q$ and $\mathbb{E}|T_k^b| < 2^k p_1q$.

Since $\{B_j\}_{j \in \mathbb{N}}$ and $\{E_j\}_{j \in \mathbb{N}}$ are independent of S_k^a (and of each other), for any $\beta > 0$ we can use the Chernoff bound (see, e.g., [6, Theorem 2.1] and [1, Appendix A]) with a deviation of $t = (\beta - 1)2^k p_2 q$ from the expectation to get

$$\mathbb{P}\left(|T_k^a| > \beta 2^k p_2 q\right) < \exp\left(-\frac{3(\beta - 1)^2}{2(\beta + 2)} 2^k p_2 q\right),$$

and analogously for Bob we have

$$\mathbb{P}\left(|T_k^b| > \beta 2^k p_1 q\right) < \exp\left(-\frac{3(\beta - 1)^2}{2(\beta + 2)} 2^k p_1 q\right).$$

Clearly, setting

$$K_a = \log_2(1/(p_2 q)) - 3, \quad K_b = \log_2(1/(p_1 q)) - 3$$

and taking β large enough (e.g., $\beta = 20$ would suffice) we get

$$\mathbb{P}\left(\bigcup_{k \geq K_a} \left\{|T_k^a| > \beta 2^k p_2 q\right\}\right) \leq 2 \mathbb{P}\left(|T_{K_a}^a| > \beta 2^{K_a} p_2 q\right) < \frac{1}{8} \quad (2.4)$$

and

$$\mathbb{P}\left(\bigcup_{k \geq K_b} \left\{|T_k^b| > \beta 2^k p_1 q\right\}\right) < \frac{1}{8}. \quad (2.5)$$

Also, since $\sum_{k < K_a} |S_k^a| < 2^{K_a} \leq (8p_2 q)^{-1}$ and similarly $\sum_{k < K_b} |S_k^b| < 2^{K_b} \leq (8p_1 q)^{-1}$, we have by Markov's inequality that

$$\mathbb{P}\left(\bigcup_{k < K_a} \left\{|T_k^a| > 0\right\}\right) \leq \sum_{k < K_a} \mathbb{E}|T_k^a| = p_2 q \sum_{k < K_a} \mathbb{E}|S_k^a| < \frac{1}{8} \quad (2.6)$$

and similarly

$$\mathbb{P}\left(\bigcup_{k < K_b} \left\{|T_k^b| > 0\right\}\right) < \frac{1}{8}. \quad (2.7)$$

Putting together (2.4),(2.5),(2.6),(2.7), with probability at least $\frac{1}{2}$ the following holds:

$$|T_k^a| \leq \begin{cases} \beta 2^k p_2 q & k \geq K_a \\ 0 & k < K_a \end{cases}, \quad |T_k^b| \leq \begin{cases} \beta 2^k p_1 q & k \geq K_b \\ 0 & k < K_b \end{cases} \quad \text{for all } k. \quad (2.8)$$

When (2.8) holds we can bound R as follows:

$$\begin{aligned} R &= \sum_j \mu_a(j) \mu_b(j) A_j B_j E_j = \sum_k \sum_\ell \sum_{j \in T_k^a \cap T_\ell^b} \mu_a(j) \mu_b(j) \leq \sum_k \sum_\ell |T_k^a \cap T_\ell^b| 2^{-k+1} 2^{-\ell+1} \\ &\leq \sum_k \sum_\ell \sqrt{|T_k^a| |T_\ell^b|} 2^{-k+1} 2^{-\ell+1} = 4 \left(\sum_k \sqrt{|T_k^a|} 2^{-k} \right) \left(\sum_\ell \sqrt{|T_\ell^b|} 2^{-\ell} \right) \\ &\leq 4\beta(p_1 p_2)^{1/2} q \left(\sum_{k \geq K_a} 2^{-k/2} \right) \left(\sum_{\ell \geq K_b} 2^{-\ell/2} \right), \end{aligned}$$

where the second inequality used the fact that $|F_1 \cap F_2| \leq \min\{|F_1|, |F_2|\} \leq \sqrt{|F_1| |F_2|}$ for any two finite sets F_1, F_2 and the last inequality applied (2.8). From here the proof is concluded by observing that

$$R \leq 16(p_1 p_2)^{1/2} q 2^{-K_a/2} 2^{-K_b/2} = 128\beta p_1 p_2 q^2. \quad \blacksquare$$

Corollary 2.3. *There exists some absolute $c > 0$ such that for any pair of stationary strategies, the expected number of rounds required for a successful synchronization is at least $c/(p_1p_2q^2)$.*

Proof. Recall that the success probability in any specific round, given the environments and strategies, is precisely R . Hence, conditioned on the value of R , the probability of synchronizing in one of the first $1/(2R)$ rounds is at most $\frac{1}{2}$. Theorem 2.2 established that with probability at least $\frac{1}{2}$ we have $R < Cp_1p_2q^2$, therefore altogether with probability at least $\frac{1}{4}$ there is no synchronizing before time $(2Cp_1p_2q^2)^{-1}$. We conclude that the statement of the corollary holds with $c = 1/(8C)$. ■

2.4 Proof of Theorem 2

Partition the channels into infinitely many sets of infinitely many channels each. On the i -th round, Alice chooses the first channel in the i -th set which is open in her environment. Bob does likewise.

Consider the probability that both parties choose the same channel: Each channel has probability $1 - (1 - p_1)(1 - p_2) = p_1 + p_2 - p_1p_2$ of being open to Alice or Bob and probability p_1p_2 of being open to both. Hence, the probability that the first channel open to either is open to both is

$$\frac{p_1p_2}{p_1 + p_2 - p_1p_2} \geq \min\{p_1, p_2\}/2.$$

If indeed both players chose the same channel at some round, it is necessarily open for both of them and with probability q it is also open in the global environment. Hence, the probability of success at each round is at least $\min\{p_1, p_2\}q/2$. Different rounds use disjoint sets of channels, so the event of success at different rounds are independent and the number of round to success has a geometric distribution with expectation $\mathbb{E}X \leq 2/(\min\{p_1, p_2\}q)$.

For a lower bound of matching order observe the following: For any strategy Alice might employ and at any given round, the probability that the channel she chooses is open for both Bob and Eve is p_2q , and this is an upper bound for the probability of success. Similar argument for Bob yields that the probability of success at any given round is at most $\min\{p_1, p_2\}q$. A straightforward first moment argument (as in the proof of Corollary 2.3) now implies that $\mathbb{E}X \geq 1/(4 \min\{p_1, p_2\}q)$.

This completes the proof. ■

References

- [1] N. Alon and J. H. Spencer, *The probabilistic method*, 3rd ed., John Wiley & Sons Inc., Hoboken, NJ, 2008.
- [2] S. Alpern, V. J. Baston, and S. Essegaier, *Rendezvous search on a graph*, J. Appl. Probab. **36** (1999), no. 1, 223–231.
- [3] S. Alpern and S. Gal, *The theory of search games and rendezvous*, International Series in Operations Research & Management Science, 55, Kluwer Academic Publishers, Boston, MA, 2003.
- [4] E. J. Anderson and R. R. Weber, *The rendezvous problem on discrete locations*, J. Appl. Probab. **27** (1990), no. 4, 839–851.

- [5] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh, *White space networking with wi-fi like connectivity*, ACM SIGCOMM Computer Communication Review **39** (2009), 27–38.
- [6] S. Janson, T. Luczak, and A. Rucinski, *Random graphs*, Wiley-Interscience Series in Discrete Mathematics and Optimization, Wiley-Interscience, New York, 2000.