Competitive Strategies for Symmetric Rendezvous on the Line*

Max Klimm† Guillaume Sagnol † Martin Skutella † Khai Van Tran †

Abstract

In the Symmetric Rendezvous Search on the Line with Unknown Initial Distance, two identical agents are placed on the real line with their distance, the other's location, and their orientation unknown to them. Moving along the line at unit speed and executing the same randomized search strategy, the agents' goal is to meet up as early as possible. The expected meeting time obviously depends on the unknown initial distance and orientations. The quality of a randomized search strategy is thus measured by its competitive ratio, that is, the ratio of the expected meeting time and the earliest possible meeting time (half the initial distance).

We present a class of successively refined randomized search strategies together with a rigorous mathematical analysis of their continuously improved competitive ratios. These strategies all rely on the basic idea of performing an infinite sequence of steps of geometrically increasing size in random directions, always returning to the agent's initial position before starting the next step. In addition, our more refined strategies use two novel ideas. First, remembering their past random choices, the agents randomly choose the direction of the next step in a Markov-chain-like manner. Second, choosing the next few random directions in advance, each agent may combine consecutive steps in the same direction into one longer step. As our main result, we show that this combination of looking into the past as well as into the future leads to a substantially improved competitive ratio of 13.93 compared to the previously best known bound of 24.85 (Ozsoyeller et al. 2013).

1 Introduction

The theory of optimal search is concerned with situations where a finite set of agents is initially placed in a given environment, and their task is to move to the same location as quickly as possible. Search strategies of this kind have applications in motion planning for teams of mobile agents or robots that need to coordinate to achieve a common task. The search problem that—according to a survey by Alpern [2], one of the founding figures of this area—received the most attention in the literature is the rendezvous problem on the line. Two agents are placed with a distance of $2d$ on an infinite line. Each agent can move at unit speed on the line. Clearly, if the agents knew their positions, they could both move towards each other and meet at time d. In rendezvous problems, however, it is assumed that the agents do not see each other and also have only a random sense of direction, i.e., at time 0 for each agent independently and uniformly at random one direction is called "forward" and the other "backward". In particular, the local sense of directions of the agents may or may not agree. How should the agents move in order to meet as quickly as possible?

For the case that the two agents know the initial distance of $2d$ and may use different strategies, Alpern and Gal [5] show that the agents meet in expectation at time $^{13d}/_{4}$, i.e., the optimal *asymmetric* search strategy is $^{13}/_{4}$-competitive. The *symmetric* rendezvous problem with known initial distance was proposed by Alpern [1] even before the asymmetric version. In symmetric rendezvous search, both agents are required to use the same (randomized) strategy. This is a sensible assumption when the agents are identical and have no other way to break ties. The use of distinct strategies by different agents presupposes that the two agents can exchange information on the type of strategy they are going to use; in the absence of possible communication, committing to symmetric strategies appears to be the only viable option. Ozsoyeller et al. also argue that the symmetric version is advantageous in robotics application as it allows for the mass manifacture of identical robots. Alpern [1] proposes a 5-competitive symmetric strategy. Despite over twenty years of research, the tight competitive ratio for symmetric rendezvous on the line with known initial distance is still unknown with the currently best lower bound being 4.152 and the currently best upper bound at 4.257 (Han et al. [12]).

Even less is known for the symmetric rendezvous problem on the line with an unknown initial distance. Baston and Gal [11] consider the more general case where the distance is drawn from an unknown probability

*Supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy — The Berlin Mathematics Research Center MATH+ (EXC-2046/1, project ID: 390685689).

†Institute of Mathematics, TU Berlin, Germany.

Copyright © 2022 by SIAM
Unauthorized reproduction of this article is prohibited

Strategy	comp. ratio
Naive Strategy (Sec. 3)	**28.86**
Baston and Gal [11].	**26.65**
Biased Strategy (Sec. 4)	**25.52**
Ozsoyeller, Beveridge, and Isler [14]	**24.85**
Markov Strategy (Sec. 5)	**20.76**
Markov Strategy with Lookahead (Sec. 6)	**13.93**

Table 1: Known and new results for symmetric rendezvous on the line with unknown distance

distribution and give a 26.65-competitive strategy. More recently, this was improved by Ozsoyeller et al. [14], who present a 24.85-competitive strategy for a deterministic but unknown initial distance.

Our results. In this paper, we present a family of new strategies for the symmetric rendezvous problem on the line with unknown distance. Our best strategy achieves competitive ratio 13.93. To achieve this significant improvement over the current state of the art, we use a couple of new ideas. Starting with a very simple strategy, we develop and analyze a series of refined strategies, employing innovative strategic features along with novel analytical techniques, finally resulting in a strategy with the competitive ratio stated above.

Our basic building block is a simple strategy with exponentially increasing step lengths. To describe the strategy, consider an agent initially placed at point 0 and suppose that the agent performs an infinite sequence of elementary operations, all starting and ending at point 0. In iteration i, the agent moves with equal probability to point α^i or to point $-\alpha^i$ for some parameter $\alpha > 1$, and then returns to the origin. A thorough analysis of this strategy reveals that it is optimal to choose $\alpha = 2/\sqrt{3}$, resulting in competitive ratio $15 + 8\sqrt{3} \approx 28.86$.

Intuitively, after not having met the other agent in iteration i, it is more likely that they were going in the wrong direction. Thus, it is sensible to increase the chance of going into the other direction in iteration $i+1$. The behavior of this strategy is naturally analyzed with a Markov chain, that can be analyzed with generating functions. Optimizing over the parameters of the strategy yields a competitive ratio of about 25.52.

We then generalize this approach allowing for longer histories, that is, the random decision taken by an agent concerning the direction of their next step may depend on the directions of several previous steps. With a generalized Markov chain approach we analyze histories up to a length of 6, which brings the competitive ratio down to 20.76, improving upon the state of the art.

Besides looking into the past, it is also worthwhile to look into the future when determining the direction of the next step. In a scenario where an agent will move into the same direction several times in a row, there is no point in always returning to the origin in-between. Replacing these steps by one longer step in the same direction increases the chance of rendezvous considerably and leads to a substantially improved competitive ratio of 13.93.

In order to increase clarity, all proofs have been moved to the appendix.

Related work. We first discuss the case of asymmetric search strategies. When the initial distance is known, in an optimal strategy one agent essentially stays in place while the other does the search. The optimal competitive ratio is 13/4, as shown by Alpern and Gal [5]. For the case that the initial distance is drawn from a probability distribution with bounded support $[0, D]$ and convex cumulative distribution function F, Alpern and Beck [3] show that a similar strategy is optimal and achieves expected rendezvous time $(4\mu + 9D)/8$, where μ is the expected distance. Optimal strategies for agents moving with different speeds are obtained by Alpern and Gal [6]. Asymmetric rendezvous problems exhibit strong connections to line search problems. A feasible strategy is that one agent stays stationary at its starting location while the other uses a line search strategy. This strategy is 2β-competitive where β is the competitiveness of the line search strategy. Furthermore, it has been argued by [5] that in the relaxation of the problem where the agents share the same orientation, it is optimal for them to always move in opposite directions, and to employ a line search strategy for this joint movement. This shows that the optimal asymmetric strategy is at best β-competitive. Using that the optimal competitive ratio for line search is 9 for deterministic strategies (Baezayates et al. [9]) and 4.591 for randomized strategies (Kao at el. [13]), this implies that the optimal competitive ratios for deterministic and randomized asymmetric search strategies are in the intervals $[9, 18]$ and $[4.591, 9.182]$, respectively. Alpern and Beck [4] give a deterministic strategy with competitive ratio 11.028 and conjecture that this is best possible.

The rendezvous problem with symmetric strategies was first considered by Alpern [1], who gives a 5-

Copyright © 2022 by SIAM
Unauthorized reproduction of this article is prohibited

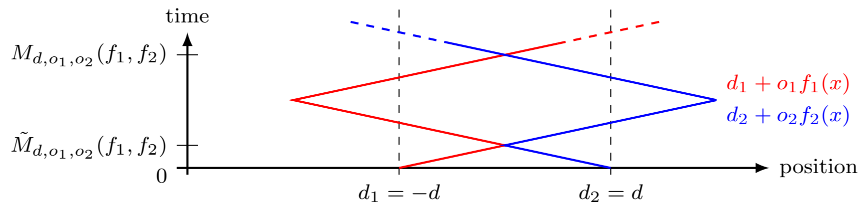

Figure 1: Movement of agent 1 (red) and agent 2 (blue) for a pair of strategies f_i and orientations o_i, $i = 1, 2$. We have $\tilde{M}_{d,o_1,o_2}(f_1, f_2) < M_{d,o_1,o_2}(f_1, f_2)$ since at time $\tilde{M}_{d,o_1,o_2}(f_1, f_2)$ both agents are at the same location, but do not cross afterwards. Note that $M_{d,o_1,o_2}(f_1, f_2)$ is robust against slight changes of d while $\tilde{M}_{d,o_1,o_2}(f_1, f_2)$ is not.

competitive randomized strategy for the case that the initial distance is known. Further improvements were obtained by Anderson and Essegaier [8] and Baston [10]. Han et al. [12] give the currently best known upper bound of 4.257 and the best known lower bound of 4.152 on the competitive ratio. For a thorough introduction into search and rendezvous search problems as well as an overview of the different settings, we refer to the textbook by Alpern and Gal [7].

2 Preliminaries

In this section we give a precise account of the Symmetric Rendezvous Problem on the line with unknown distance. We consider the following setting. An adversary chooses a value $d \in \mathbb{R}_{>0}$ and two orientations $o_1, o_2 \in \{-1, 1\}$[1]. Two identical agents are placed at points $d_1 := -d$ and $d_2 := d$ of an infinite line. The agents do not know the choices made by the adversary.

A deterministic strategy (or pure strategy) of an agent is a Lipschitz continuous function $f : \mathbb{R}_{\geq 0} \to \mathbb{R}$ with Lipschitz constant 1, i.e., $|f(x) - f(x')| \leq |x - x'|$ for all $x, x' \in \mathbb{R}_{\geq 0}$. The value $f(x)$ is the position of the agent at time x relative to its start position and relative to the initial orientation of the agent, which is chosen by the adversary. More specifically, if agent i follows strategy f and its initial orientation is o_i, then its position at time x is $d_i + o_i f(x)$ for all $x \geq 0$. A randomized strategy $\mathcal{S}$ is a probability distribution over deterministic strategies.

In the following, we restrict to the case of *symmetric strategies*, that is, both agents must follow the same deterministic or randomized strategy. In this symmetric setting, however, it is easy to see that for every deterministic strategy f, the agents will always remain at distance $2d$ when $o_1 = o_2$. So in order to ensure that the agents can eventually meet, randomization is essential. For a pair of deterministic strategies f_1, f_2, we define their *meeting time* as

$$M_{d,o_1,o_2}(f_1, f_2) \;:=\; \inf\{x \in \mathbb{R}_{\geq 0} : d_1 + o_1 f_1(x) > d_2 + o_2 f_2(x)\}.$$

That is, the meeting time is the first point in time when the agents are about to cross. The competitive ratio of a randomized strategy $\mathcal{S}$ is

$$\rho_{\mathcal{S}} \;:=\; \sup_{d \in \mathbb{R}_{>0}, o_1, o_2 \in \{-1,1\}} \frac{\mathbb{E}_{f_1, f_2 \sim \mathcal{S}}\left[M_{d,o_1,o_2}(f_1, f_2)\right]}{d}. \tag{2.1}$$

We note that usually the meeting time is defined as the first point in time at which both agents are at the same position, i.e., $\tilde{M}_{d,o_1,o_2}(f_1, f_2) := \min\{x \in \mathbb{R}_{\geq 0} : d_1 + o_1 f_1(x) = d_2 + o_2 f_2(x)\}$. The latter definition, however, has the disadvantage that, even for very natural strategies, the supremum over d in (2.1) might not be attained; see Figure 1 for an illustration.

The randomized strategies $\mathcal{S}$ that we devise and analyze in this paper have the property that for every $f \sim \mathcal{S}$ and for every $\varepsilon > 0$, there are $x, y \in [0, \varepsilon]$ such that $f(x) < 0$ and $f(y) > 0$, i.e., for every realization $f \sim \mathcal{S}$, the pure strategy f oscillates around 0 arbitrarily close to the starting time 0. While this behavior is consistent with our definition of a strategy, such strategies may not be implementable, e.g., by a Turing machine, since there is no clear first step after time 0. A couple of remarks are in order.

[1]This setting with adversarial orientations is clearly harder than the more common case where they are chosen at random; thus, the results presented in this paper also hold for orientations o_1, o_2 resulting from a coin flip.

Copyright © 2022 by SIAM
Unauthorized reproduction of this article is prohibited

First, it is not hard to see that some sort of oscillation of this kind is necessary in order to achieve a constant competitive ratio. To see this, suppose that there is $\varepsilon > 0$ such that

$$\mathbb{P}_{f\sim\mathcal{S}}\big[f(x) \geq 0 \text{ for all } x \in [0,\varepsilon]\big] > 0.$$

The case where $\mathbb{P}_{f\sim\mathcal{S}}\big[f(x) \leq 0 \text{ for all } x \in [0,\varepsilon]\big] > 0$ is analogous and, thus, omitted. Let $k \in \mathbb{Z}_{>0}$ be arbitrary, and consider the adversarial choice $d = \varepsilon/k$ together with $o_1 = -1$ and $o_2 = 1$. Let $p := \mathbb{P}_{f\sim\mathcal{S}}\big[f(x) \geq 0 \text{ for all } x \in [0,\varepsilon]\big] > 0$. We then obtain $\mathbb{E}_{f_1,f_2\sim\mathcal{S}}\big[M_{d,o_1,o_2}(f_1,f_2)\big] > p^2\varepsilon$, since with probability at least p^2, agent 1 is always left of $d_1 = -d < 0$ and agent 2 is always right of $d_2 = d > 0$ within the time interval $[0,\varepsilon]$. This implies $\rho_{\mathcal{S}} > kp^2$. As k was chosen arbitrarily, the competitive ratio is thus unbounded. Second, similar oscillating behavior has been used before in the literature, e.g., by Alpern and Beck [3] as well as Baston and Gal [11]. Third, note that the oscillation arbitrarily close to zero may be avoided if we assume that the adversary can only choose $d \in [d_0,\infty]$ for some arbitrary constant $d_0 > 0$. This allows to skip over some time interval $[0,\varepsilon]$ for some $\varepsilon > 0$ where rendezvous cannot occur, only improving the competitive ratio. With this discussion in mind, we define the following set of strategies.

Definition 1. *Let $(x_i,y_i)_{i\in\mathbb{Z}}$ be a sequence with $x_i \in \mathbb{R}_{>0}$ and $y_i \in \mathbb{R}$ for all $i \in \mathbb{Z}$ such that the following properties are satisfied: (i) $x_i < x_{i+1}$ for all $i \in \mathbb{Z}$; (ii) $x_{i+1} - x_i = |y_{i+1} - y_i|$ for all $i \in \mathbb{Z}$; (iii) $\lim_{i\to-\infty} x_i = \lim_{i\to-\infty} y_i = 0$ and $\lim_{i\to\infty} x_i = \infty$. The corresponding deterministic strategy is defined as $f(x) := y_i + \frac{y_{i+1}-y_i}{x_{i+1}-x_i}(x - x_i)$ for $x \in (x_i, x_{i+1}]$, $i \in \mathbb{Z}$, and $f(0) = 0$.*

It is straightforward to verify that the deterministic strategy f is well-defined. Indeed, since $\lim_{i\to-\infty} x_i = 0$, $\lim_{i\to\infty} x_i = \infty$ and $x_i < x_{i+1}$, for every $x > 0$ there is a unique $i \in \mathbb{Z}$ such that $x \in (x_i, x_{i+1}]$. To see that f is Lipschitz continuous with constant 1, note that f is continuous and differentiable on $\mathbb{R}_{>0} \setminus \{x_i \mid i \in \mathbb{Z}\}$ and thus differentiable almost everywhere. For $x \in (x_i, x_{i+1})$ for some $i \in \mathbb{Z}$, we obtain $|f'(x)| = \big|(y_{i+1} - y_i)/(x_{i+1} - x_i)\big| = 1$, and the result follows.

For ease of exposition, in the following, we often specify a deterministic strategy f only at the points x_i for a sequence as in Definition 1, i.e., setting $f(x_i) := y_i$, with the understanding that f is then completed as in Definition 1. Similarly, we specify random strategies by giving a probability distribution of sequences $(x_i,y_i)_{i\in\mathbb{Z}}$ as in Definition 1, which thus implies a probability distribution of corresponding deterministic strategies.

3 Naive Strategy

In this section we introduce and analyze the Naive Strategy $\mathsf{NS}(\alpha)$, that lays the basic framework for all strategies in this paper. The strategy is described relative to a given scaling factor $\alpha > 1$ to be optimized later.

Naive Strategy $\mathsf{NS}(\alpha)$

1. Sample $(c_i)_{i\in\mathbb{Z}} \in \{-1,1\}^{\mathbb{Z}}$ such that $\mathbb{P}[c_i = +1] = 1/2$ independently for all $i \in \mathbb{Z}$.
2. For $i \in \mathbb{Z}$, let $x_{2i} := \sum_{k=-\infty}^{i-1} 2\alpha^k$, $y_{2i} := 0$ and $x_{2i+1} := x_{2i} + \alpha^i$, $y_{2i+1} := c_i\alpha^i$.
3. Use randomized strategy with time points $(x_i)_{i\in\mathbb{Z}}$ and random positions $(y_i)_{i\in\mathbb{Z}}$.

We subdivide time into consecutive iterations $i \in \mathbb{Z}$, where iteration i takes $2\alpha^i$ time. The agents start each iteration at their initial starting position. In iteration i, each agent randomly makes a step of length α^i. The step is taken with probability $1/2$ into the positive direction to the point α^i and with probability $1/2$ in the negative direction to the point $-\alpha^i$. After reaching one of these points, the agent immediately returns to its starting location to begin the next iteration. Thus, iteration i starts at the initial position $y_{2i} := 0$ at time $x_{2i} := \sum_{k=-\infty}^{i-1} 2\alpha^k$, after all preceding iterations are completed. Moving to $y_{2i+1} := \pm\alpha^i$ takes α^i time such that $x_{2i+1} := x_{2i} + \alpha^i$. An illustration is given in Figure 2.

Theorem 3.1. *If $1 < \alpha < 4/3$, the competitive ratio of $\mathsf{NS}(\alpha)$ is $\rho_{\mathsf{NS}(\alpha)} = 1 - \frac{2\alpha}{3\alpha^2-7\alpha+4}$.*

Optimizing this ratio over $\alpha \in (1, 4/3)$, we obtain the following direct corollary of Theorem 3.1.

Corollary 3.1. *The best possible competitive ratio for the* Naive Strategy *is attained for $\alpha = 2/\sqrt{3} \approx 1.16$ and yields competitive ratio $\rho_{\mathsf{NS}(\alpha)} = 15 + 8\sqrt{3} < 28.857$.*

Copyright © 2022 by SIAM
Unauthorized reproduction of this article is prohibited

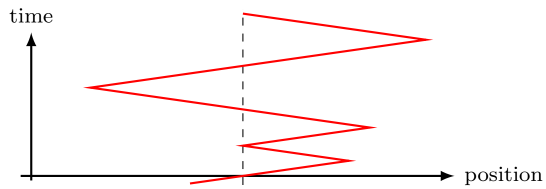

Figure 2: The dashed line represents the starting location of an agent. Displayed in red is part of the trace of an agent executing the NAIVE STRATEGY. The four displayed iterations correspond to the excerpt $(\cdots, +1, +1, -1, +1, \cdots)$ from the series of choices $(c_i)_{i\in\mathbb{Z}}$.

4 Biased Strategy

In this and the next section we introduce iterative refinements of the NAIVE STRATEGY. We retain the basic movement framework but make smarter decisions when it comes to choosing the next side to go towards. That is, we generate a more elaborate series of random choices $(c_i)_{i\in\mathbb{Z}}$.

In iteration i of the NAIVE STRATEGY, each agent decides to make a move from its origin to $-\alpha^i$ or α^i via a fair coin flip. If the agents do not meet in an iteration i, there are two possible reasons for that: The step size might be too small, that is, smaller than half the starting distance (i.e., $\alpha^i < d$); in this case, the agents can do nothing but wait for an iteration with sufficiently large step size. Otherwise at least one agent has chosen the wrong direction; in that case, the agent is incentivized to choose the other direction in iteration $i+1$ with a probability p that is larger than $1/2$. This is the main intuition behind the BIASED STRATEGY $\mathsf{BS}(\alpha, p)$, that depends both on the scaling parameter α and the probability p that the direction is changed in the next iteration. Both parameters will be optimized later.

BIASED STRATEGY $\mathsf{BS}(\alpha, p)$

1. Sample $(c_i)_{i\in\mathbb{Z}} \in \{-1,1\}^{\mathbb{Z}}$ such that $\mathbb{P}[c_{i+1} = -c_i] = p$ independently for all $i \in \mathbb{Z}$.
2. Continue as in the NAIVE STRATEGY $\mathsf{NS}(\alpha)$.

In the first step of the strategy, we sample a sequence $(c_i)_{i\in\mathbb{Z}} \in \{-1, 1\}$ with the property that $\mathbb{P}[c_{i+1} = -c_i] = p$ for all $i \in \mathbb{Z}$. This can be done as follows. We first set $c_0 = \pm 1$ by a fair coin flip. We then iteratively set $c_{i+1} = -c_i$ with probability p and $c_{i+1} = c_i$ with the remaining probability for all $i \in \mathbb{Z}_{\geq 0}$. Then we iteratively set $c_{i-1} = -c_i$ with probability p and $c_{i-1} = c_i$ with the remaining probability for all $i \in \mathbb{Z}_{<0}$. Notice that by this construction, since $c_0 = -1$ and $c_0 = 1$ appear with equal probability, we have that

$$\mathbb{P}[c_i = 1] = \mathbb{P}[c_i = -1] = \tfrac{1}{2} \quad \text{for all } i \in \mathbb{Z}. \tag{4.2}$$

THEOREM 4.1. *Using the* BIASED STRATEGY *with* $\alpha = 1.179$ *and* $p = 0.679$ *we can achieve competitive ratio* $\rho_{\mathsf{BS}(\alpha,p)} < 25.513$.

To prove this result, we first introduce random variables U_i, V_i, and W_i that indicate whether the agents run towards each other, in parallel, or move away from each other in the ith iteration. The result follows by establishing a recurrence equation linking the probabilities of these events between successive iterations, and by using the method of generating functions, cf. Appendix.

5 Markov Strategy

The BIASED STRATEGY achieves an improvement over the NAIVE STRATEGY by exploiting the fact that the previous decision did not result in rendezvous. A natural generalization is to exploit the fact that the h previous decisions did not result in rendezvous for some $h \geq 2$.

Extending the intuition behind the BIASED STRATEGY, it seems reasonable to assign a higher switching probability to an agent that has visited the same side several times in a row, whereas an agent that has switched sides every time should rather stay on one side with a higher probability. Formally, we choose a memory depth of h bits and encode the previous h choices as bits, where a 0 denotes choosing the same side as in the previous iteration and a 1 denotes choosing the opposite side. This yields a state space of $\{0,1\}^h$, where the agent is in exactly one state in each iteration.

Copyright © 2022 by SIAM
Unauthorized reproduction of this article is prohibited

Denote by $\sigma \in \{0,1\}^h$ the state of an agent in iteration i. Then, the agent changes direction in iteration $i+1$ with probability p_σ, taking them to the state $(\sigma_2, \dots, \sigma_h, 1)$. With probability $1 - p_\sigma$, the agent sticks to their previous direction in iteration $i+1$, taking them to state $(\sigma_2, \dots, \sigma_h, 0)$.

Markov Strategy: $\mathsf{MS}(\alpha, h, p)$

1. Choose $a \in \{0,1\}$ uniformly at random and sample a bit-sequence $(a_i)_{i\in\mathbb{Z}} \in \{0,1\}^{\mathbb{Z}}$ from the probability distribution described below, i.e., $\mathbb{P}[a_i = 1] = p_{(a_{i-h},\cdots,a_{i-1})}$ for all $i \in \mathbb{Z}$.
2. Set $c_i := \begin{cases} (-1)^{a+\sum_{k=i+1}^{0} a_k} & \text{for} \quad i < 0, \\ (-1)^{a+\sum_{k=1}^{i} a_k} & \text{for} \quad i \geq 0. \end{cases}$
3. Continue as in the Naive Strategy $\mathsf{NS}(\alpha)$.

To summarize, the agent's dynamics is fully described by a discrete time Markov chain (which gives rise to the name of the strategy) with transition matrix $\underline{A}$, whose entries are given by

$$\underline{A}_{\sigma,\sigma'} := \begin{cases} p_\sigma & \text{if} \quad \sigma' = (\sigma_2, \dots, \sigma_h, 1), \\ 1 - p_\sigma & \text{if} \quad \sigma' = (\sigma_2, \dots, \sigma_h, 0), \\ 0 & \text{otherwise.} \end{cases}$$

Note that this Markov chain may have some transient states: for example, if we always change direction after three iterations in the same direction, then all states containing three consecutive 0-bits are transient. Without loss of generality, we can restrict the Markov chain to a subset of recurrent states, since the other states occur with zero probability after the initial trembling.

If the above-defined Markov chain has a unique stationary distribution, we can sample the sequence $(a_i)_{i\in\mathbb{Z}}$ in the first step of the strategy as follows. Let s_σ be the probability of being in state σ in the stationary distribution. First, we sample $(a_0, \cdots, a_{h-1})$ according to the unique stationary distribution. We can then iteratively sample a_i for $i \geq h$ according to the Markov chain, i.e. $\mathbb{P}[a_i = 1] = p_{(a_{i-h},\cdots,a_{i-1})}$ for $i \geq h$. As for the iterations $i < 0$, we first note that in general a Markov chain is not time-reversible (we cannot assume our Markov chain to fulfill the detailed balance equations, which are a necessary condition for general time reversibility). However we know that $(a_{i-h+1}, \cdots, a_i)$ is distributed according to the unique stationary distribution of the Markov chain, since there is an infinite number of preceding Markov transitions. This is sufficient to ensure time-reversibility and we can sample a_i for $i < 0$ according to Bayes. Let $i < 0$, $\bar{\sigma} = (a_{i+1}, \cdots, a_{i+h-1})$, and a_{i+h} be fixed, then $\mathbb{P}[a_i = 1] = \underline{A}_{(1,\bar{\sigma}),(\bar{\sigma},a_{i+h})} \cdot s_{(1,\bar{\sigma})} / s_{(\bar{\sigma},a_{i+h})}$.

In order to analyze this strategy, we track another Markov chain that stores the *actual directions* taken by both agents (rather than just their switching sequence). We first describe this chain for a single agent. The *actual state* of an agent is denoted by $\boldsymbol{s} \in \{\texttt{True}, \texttt{False}\}^{h+1}$. The first h bits indicate whether the agent was moving towards (`True`) or away from (`False`) the other agent during the previous h iterations; the last bit s_{h+1} indicates the direction chosen for the current iteration.

Denote by $\sigma(\boldsymbol{s}) \in \{0,1\}^h$ the sequence of switches corresponding to $\boldsymbol{s} \in \{\texttt{True}, \texttt{False}\}^{h+1}$, i.e.,

$$\sigma(\boldsymbol{s})_i := \begin{cases} 0 & \text{if } s_i = s_{i+1}, \\ 1 & \text{otherwise,} \end{cases} \qquad \text{for } i = 1, \dots, h.$$

Similarly as before, the transition matrix A between the actual states can be defined as follows:

$$A_{\boldsymbol{s},\boldsymbol{s}'} = \begin{cases} p_{\sigma(\boldsymbol{s})} & \text{if} \quad \boldsymbol{s}' = (s_2, \dots, s_{h+1}, \neg s_{h+1}), \\ 1 - p_{\sigma(\boldsymbol{s})} & \text{if} \quad \boldsymbol{s}' = (s_2, \dots, s_{h+1}, s_{h+1}), \\ 0 & \text{otherwise.} \end{cases}$$

As before, we restrict attention to the recurrent states $\mathcal{R} \subseteq \{\texttt{True}, \texttt{False}\}^{h+1}$ of this Markov chain. Denote by $\pi_i \in [0,1]^{|\mathcal{R}|}$ the vector of probabilities for being in a certain state in iteration i. By definition, we have $\pi_{i+1} = \pi_i A$. If there is a unique stationary distribution, we denote it by π.

Copyright © 2022 by SIAM
Unauthorized reproduction of this article is prohibited

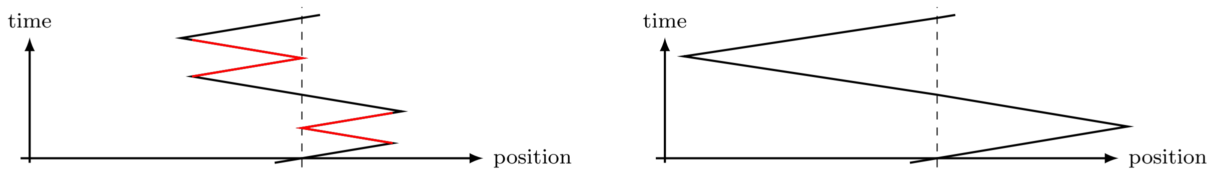

Figure 3: On the left hand side we have the trace of an agent going to the same side twice in a row, where the dashed line represents the starting location. It is impossible for rendezvous to occur on the red part of the trace. Thus, "bending" the path outwards and combining both iterations into a single "lunge" as displayed on the right hand side, cannot increase the competitive ratio.

Finally, we introduce a Markov chain tracking the entire execution of our strategy. We want to track the combined state $(\boldsymbol{s}^1, \boldsymbol{s}^2)$ of both agents in a Markov chain. As long as $\alpha^i \leq d$, the agents cannot meet and this Markov chain is basically the Kronecker product of the previous Markov chain with itself, so its transition matrix is $\bar{A} = A \otimes A$. Moreover, if there is a unique stationary distribution, it is $\bar{\pi} = \pi \otimes \pi$.

Let $i_0 := \lfloor 1 + \log_\alpha d \rfloor$ denote the first iteration in which the agents can meet. From iteration i_0 on, rendezvous occurs if and only if both agents choose the correct side. We denote by S_0 the corresponding subset of combined states $S_0 := \{(\boldsymbol{s}^1, \boldsymbol{s}^2) \in \mathcal{R}^2 : s^1_{h+1} = s^2_{h+1} = \texttt{True}\}$, and we let $S_+ := (\mathcal{R} \times \mathcal{R}) \setminus S_0$ denote the remaining states.

Theorem 5.1. *If the switching probabilities (p_σ) and α are chosen such that (i) the Markov chain with transition matrix A has a unique stationary distribution π; and (ii) $\alpha \bar{A}_+$ has a spectral radius less than 1, then the competitive ratio of the* Markov Strategy *is*

$$\rho_{\mathsf{MS}(\alpha,h,p)} = \frac{2\alpha}{\alpha - 1} + 1 + 2\alpha\, \bar{\pi}_+ \left(I - \alpha \bar{A}_+\right)^{-1} \mathbf{1},$$

where $\bar{\pi}_+ = \bar{\pi}[S_+]$ is the vector with the entries of $\bar{\pi} = \pi \otimes \pi$ indexed in S_+, $\bar{A}_+ = \bar{A}[S_+, S_+]$ is the principal submatrix of $\bar{A} = A \otimes A$ corresponding to rows and columns indexed in S_+, $\mathbf{1}$ is the vector of all-ones, and I denotes the identity matrix.

The idea behind this result is that the time spent before the first iteration in which the agents can meet contributes $2\alpha(\alpha - 1)^{-1}$ to the competitive ratio; the iteration in which the agents actually meet contributes an additional 1; the last term follows from an analysis of the time spent in the remaining iterations and involves a matrix power series, that converges to $\left(I - \alpha \bar{A}_+\right)^{-1}$, cf. Appendix.

The task of optimizing $\rho_{\mathsf{MS}}(\alpha, h, p)$ with respect to p and α is a nonconvex and nondifferentiable optimization problem. We used a simple black-box optimization routine to find the best possible competitive ratios. The best found local minima for different values of h are reported in the first column of Table 2; in particular, for a history length of $h = 6$ bits we get the following.

Corollary 5.1. *Using a history length of $h = 6$ bits, $\alpha = 1.238$ and the p_σ's defined in Table 3 (in the appendix), a competitive ratio of $\rho_{\mathsf{MS}(\alpha,h,p)} < 20.7533$ can be achieved.*

6 Markov Strategy with Lookahead

The final improvement presented in this paper yields the Markov Strategy with Lookahead: In iteration i an agent moves out to $\pm\alpha^i$ hoping to achieve rendezvous. If this does not happen, it moves back to its starting position. Since the other agent is also moving back to its starting position, it is impossible to achieve rendezvous on the way back. Assume the agent picks the same side again in iteration $i + 1$ and moves to $\pm\alpha^{i+1}$. We can only achieve rendezvous if both agents choose the correct side, thus it is only possible to achieve rendezvous in iteration $i + 1$ if the agent already chose the correct side in iteration i. Thus, instead of moving from $\pm\alpha^i$ to 0 and then back via $\pm\alpha^i$ towards $\pm\alpha^{i+1}$, a wiser strategy looks ahead and combines both iterations into a single "lunge" from 0 to $\pm(\alpha^i + \alpha^{i+1})$, only arriving back at 0 at the end of iteration $i + 1$; see Figure 3.

Taking this idea one step further, we might want to look ahead several iterations and combine multiple consecutive visits to the same side into a single bigger "lunge". Note that from a prescriptive point of view, during a lunge the agent has to "look ahead" several iterations to know when it is time to turn back towards the origin (though note that with the restriction below the robot only ever has to look ahead L iterations). From

Copyright © 2022 by SIAM
Unauthorized reproduction of this article is prohibited

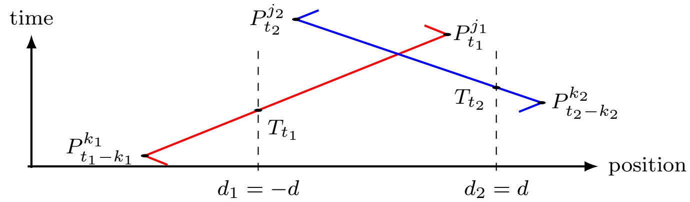

Figure 4: The dashed lines represent the starting positions of the two agents. The thick lines are the paths of two agents that have made a lunge towards the wrong side of sizes k_1 and k_2, respectively, before making a lunge in the correct direction of sizes j_1 and j_2, respectively. The latter lunges start in iteration t_1 and t_2, respectively. We have annotated the times at which the peaks occur as well as the start of iterations t_1 and t_2.

a descriptive point of view, all (random) choices can be fixed prior to any movement happening anyways. We will also refer to a movement that only spans one iteration as a lunge. To simplify the analysis, we will restrict ourselves to a maximum lookahead of size $L \in \mathbb{Z}_{>0}$. In practice, this means that we change direction after any bit-sequences σ that end with $L-1$ zeros. We thus assume that

$$\sigma_{h-L+2} = 0 \wedge \ldots \wedge \sigma_h = 0 \quad \Longrightarrow \quad p_\sigma = 1. \tag{6.3}$$

As described above, this restriction intuitively makes sense and the best parameters for the Markov Strategy we were ably to identify comply with the restriction for $L = 4$, see Table 3 (in the appendix). For the Markov Strategy with Lookahead our best found parameters yield significantly diminishing returns for an increase in L for a fixed history length h, see Table 2.

Markov Strategy with Lookahead: $\mathsf{MSL}(\alpha, h, L, p)$

1. Choose $a \in \{0, 1\}$ uniformly at random and sample a bit-sequence $(a_i)_{i\in\mathbb{Z}} \in \{0,1\}^{\mathbb{Z}}$ from the probability distribution described above, i.e., $\mathbb{P}[a_i = 1] = p_{(a_{i-h},\cdots,a_{i-1})}$ for all $i \in \mathbb{Z}$.
2. Decompose $(a_i)_{i\in\mathbb{Z}}$ into blocks of length in $\{1, 2, \cdots, L\}$ and of the form $10^\star$ corresponding to one lunge, respectively, i.e., consider the sequence $(t_j, k_j)_{j\in\mathbb{Z}}$ where $a_{t_j} = 1$, $t_{j+1} = t_j + k_j$, $a_{t_j+\ell} = 0$ for all $1 \le \ell < k_j \le L, j \in \mathbb{Z}$, and $t_0 = \min\{i : i \ge 0 \wedge t_i = 1\}$.
3. For $j \in \mathbb{Z}$, let $x_{2j} := T_{t_j} := \sum_{p=-\infty}^{t_j-1} 2\alpha^p$, $y_{2j} := 0$ and $x_{2j+1} := P_{t_j}^{k_j} := x_{2j} + \sum_{p=t_j}^{t_j+k_j-1} \alpha^p$, $y_{2j+1} := (-1)^{a+j} \sum_{p=t_j}^{t_j+k_j-1} \alpha^p$.
4. Use randomized strategy with time points $(x_j)_{j\in\mathbb{Z}}$ and random positions $(y_j)_{j\in\mathbb{Z}}$.

All prior strategies have a unique (up to scaling by α) worst case starting distance and both agents move exactly d in the iteration in which rendezvous is achieved. Due to the possible lunges (of different sizes) this is no longer the case in the Markov Strategy with Lookahead. Let $T_i := \sum_{p=-\infty}^{i-1} 2\alpha^p = \alpha^i \frac{2}{\alpha-1}$ be the starting time of iteration i. Furthermore, let $\ell_i^k := \alpha^i(1 + \ldots + \alpha^{k-1}) = \alpha^i \cdot \frac{\alpha^k-1}{\alpha-1}$ be the size of a lunge starting in iteration i and spanning k iterations in total. Moreover, let $P_i^k := T_i + \ell_i^k = \alpha^i \frac{\alpha^k+1}{\alpha-1}$ be the "peak time" of that lunge. We slightly abuse notation and define $\ell_i^k := \alpha^i \cdot \frac{\alpha^k-1}{\alpha-1}$ for any $k \in \mathbb{Z}$. So we have $\ell_i^0 = 0$ and $\ell_i^k < 0$ for $k < 0$.

Lemma 6.1. *Rendezvous can only be achieved at times $P_i^\delta + d$ with $0 \le \delta \le L-1$. Moreover, if rendezvous occurs at time $P_i^\delta + d$, then the two agents must have started lunges in the correct direction at times T_i and $T_{i+\delta}$, respectively.*

Lemma 6.2. *Assume that the agents are starting, at iterations t_1 and t_2, lunges of sizes j_1 and j_2 towards the other agent, respectively, with $t_1 \le t_2$. Assuming that no rendezvous occurs at any time $P_i^\delta + d$ with $(i, \delta) < (t_1, t_2 - t_1)$ —in the lexicographic order—, then rendezvous occurs at time $P_{t_1}^{t_2-t_1} + d$ if and only if $\min(\ell_{t_2}^{t_1+j_1-t_2}, \ell_{t_1}^{t_2+j_2-t_1}) > d$.*

Corollary 6.1. *If the strategy achieves rendezvous with probability 1, then the worst case starting distance is one of ℓ_i^k, with $1 \le k \le L$ and $i \in \mathbb{Z}$.*

Copyright © 2022 by SIAM
Unauthorized reproduction of this article is prohibited

As in the analysis of the Markov Strategy, we build a Markov chain that tracks the entire execution of our strategy. In order to do so, we interpret a state $\boldsymbol{s} \in \{\texttt{True}, \texttt{False}\}^{h+1}$ with $h \geq L$ as follows: For simplicity, we denote the coordinates of a state $\boldsymbol{s} \in \{\texttt{True}, \texttt{False}\}^{h+1}$ by $(s_{-(h-L+1)}, \ldots, s_{-1}, s_0, \ldots, s_{L-1})$. Then, $s_k = \texttt{True}$ ($s_k = \texttt{False}$) in iteration i means that the agent is moving towards (away from) the other agent in iteration $i+k$. Note that the state now stores information about the next L choices of the agents, in order to know the size of the current lunge, and to account for the other agent starting a lunge several iterations later that leads to rendezvous.

As in the previous section, we restrict to a Markov chain of recurrent states $\mathcal{R} \subseteq \{\texttt{True}, \texttt{False}\}^{h+1}$ and track the combined states $(\boldsymbol{s}^1, \boldsymbol{s}^2) \in \mathcal{R}^2$. As long as rendezvous is impossible, the transition matrix is the same as in the previous section: $\bar{A} = A \otimes A$. Things get more complicated as soon as rendezvous may start to occur. In the analysis we assign a rendezvous occurring at time $P_i^\delta + d$ to iteration i. We partition the states of iteration i. Consider a combined state $(\boldsymbol{s}^1, \boldsymbol{s}^2) \in \mathcal{R}^2$ in iteration i, and assume no rendezvous assigned to an iteration $i' < i$ has occurred. Then, we write

- $(\boldsymbol{s}^1, \boldsymbol{s}^2) \in S_i^\delta(d)$ if rendezvous occurs at time $P_i^\delta + d$, for some $0 \leq \delta \leq L-1$;
- $(\boldsymbol{s}^1, \boldsymbol{s}^2) \in S_i^{+}(d)$ if no rendezvous assigned to iteration i occurs.

By construction, for all $i \in \mathbb{Z}$ we have the partition $\mathcal{R}^2 = \left(\bigcup_{\delta=0}^{L-1} S_i^\delta(d)\right) \cup S_i^{+}(d)$.

Lemma 6.3. *For all $(\boldsymbol{s}^1, \boldsymbol{s}^2) \in \mathcal{R}^2$ and for all $i \in \mathbb{Z}$, we can efficiently decide the class $\mathsf{c} \in \{0, \ldots, L-1, +\}$ such that $(\boldsymbol{s}^1, \boldsymbol{s}^2) \in S_i^{\mathsf{c}}(d)$.*

We also partition the iterations of the strategy into 3 continuous regimes: The first regime contains all iterations in which rendezvous cannot occur. The intermediate regime starts with iteration i_0 and contains all iterations i where the classification of states into $S_i^*(d)$ depends on the specific iteration i they are in. The final regime starts with iteration i_1, where i_1 is the first iteration such that the partition of $\mathcal{R}^2$ into the $S_i^*(d)$'s is the same for all $i \geq i_1$.

Proposition 6.1. *Using Lemma 6.1 we can determine the boundaries of the regimes:*
(i) The first iteration at which rendezvous can occur is $i_0 = \lfloor \log_\alpha d/\ell_0^L \rfloor + 1$
(ii) We enter the final regime at iteration $i_1 = \lfloor \log_\alpha d \rfloor + 1$.

Theorem 6.1. *For a fixed initial distance d, denote by $\bar{A}_i(d)$ the transition matrix at iteration i, i.e., $\bar{A}_i(d)$ is the matrix $\bar{A} = A \otimes A$ with rows indexed in $S_i^k(d)$ zeroed out for $0 \leq k \leq L-1$. If the switching probabilities (p_σ) and parameter α are chosen such that (i) the Markov Chain with transition matrix A has a unique stationary distribution π; and (ii) the matrix $\alpha \bar{A}_{i_1}(d)$ has spectral radius less than 1 for each $d \in \{\ell_0^1, \ldots, \ell_0^L\}$, then the competitive ratio of the* Markov Strategy with Lookahead *is*

$$\rho_{\mathsf{MSL}(\alpha,h,L,p)} = \max_{d \in \{\ell_0^1, \cdots, \ell_0^L\}} \quad 1 + \frac{1}{d} \left\{ \sum_{i=i_0}^{i_1-1} \alpha^i \sum_{\delta=0}^{L-1} \frac{\alpha^\delta + 1}{\alpha - 1} \cdot \bar{\pi}_i \cdot \boldsymbol{\chi}_{S_i^\delta(d)} \right.$$
$$\left. + \alpha^{i_1} \cdot \bar{\pi}_{i_1} \cdot (I - \alpha \bar{A}_{i_1}(d))^{-1} \cdot \left(\sum_{\delta=0}^{L-1} \frac{\alpha^\delta + 1}{\alpha - 1} \cdot \boldsymbol{\chi}_{S_{i_1}^\delta(d)} \right) \right\},$$

where $\bar{\pi}_{i_0} := \pi \otimes \pi$, $\bar{\pi}_i := \bar{A}_i(d)^{i-i_0}$ for all $i \geq i_0$, and $\boldsymbol{\chi}_S \in \{0,1\}^{\mathcal{R}^2}$ denote the incidence vector of the subset $S \subseteq \mathcal{R}^2$.

Remark 6.1. *In the above expression for ρ_{MSL}, the indices $i_0 = i_0(d)$, $i_1 = i_1(d)$ and the vectors $\bar{\pi}_i = \bar{\pi}_i(d)$ actually depend on the value of $d \in \{\ell_0^1, \cdots, \ell_0^L\}$; we have not indicated this dependency for the sake of readability.*

The proof of this result is analog to the proof of Theorem 5.1, but contains an additional analysis for the iterations of the intermediate regime. Moreover, we must pay attention to the fact that several worst-case distances can occur, cf. Appendix.

We tried to find the best possible switching probabilities (p_σ) and scaling factor α by using a blackbox optimization routine. This is a hard, nonconvex and nondifferentiable optimization problem, with a single evaluation of $\rho_{\mathsf{MSL}}(\alpha, h, L, p)$ already requiring several seconds for a history depth of $h = 6$, thus we do not

Copyright © 2022 by SIAM
Unauthorized reproduction of this article is prohibited

claim to have solved it to optimality. Nevertheless, the best found competitive ratios are indicated in Table 2, for different values of h and L. The table clearly shows the power of the look-ahead strategies compared to the standard Markov stragegy from the previous section. In particular, allowing lunges of maximum size $L = 2$ already yields an improvement over the best found Markov Strategy without lookahead, even if we restrict the history depth to $h = 2$. By setting $p_{01} = 0.434$, $p_{10} = 1$, $p_{11} = 0.686$, and $\alpha = 1.21418$ we obtain a strategy with a very compact representation having competitive ratio $\rho_{\mathsf{MSL}(\alpha,h,L,p)} < 20.3070$. The best strategy found overall for $h = L = 6$ gives the following corollary; its competitive ratio is plotted against the unknown value $d \in [1, \alpha)$ in Figure 5 (in the appendix).

COROLLARY 6.2. *Using the* MARKOV STRATEGY WITH LOOKAHEAD *with* $\alpha = 1.18684$, $h = 6$, $L = 6$, *and the switching probabilities* p *depicted in Table 4 (in the appendix), a competitive ratio of* $\rho_{\mathsf{MSL}(\alpha,h,L,p)} < 13.9262$ *can be achieved.*

	no lookahead		$L = 2$		$L = 3$		$L = 4$		$L = 5$		$L = 6$	
h	ρ_{MS}	α	ρ_{MSL}	α	ρ_{MSL}	α	ρ_{MSL}	α	ρ_{MSL}	α	ρ_{MSL}	α
0	25.52	1.178										
1	23.52	1.197										
2	22.36	1.211	20.31	1.214								
3	21.54	1.223	19.87	1.214	16.49	1.200						
4	21.11	1.230	19.65	1.214	16.30	1.200	15.04	1.191				
5	20.80	1.237	19.48	1.214	16.02	1.200	14.68	1.188	14.20	1.187		
6	20.76	1.237	19.42	1.214	15.88	1.198	14.60	1.189	13.99	1.187	13.93	1.187

Table 2: Best found competitive ratios (rounded up) for the Markov Strategy without lookahead (ρ_{MS}) or with lookahead (ρ_{MSL}), for different values of the maximum lunge size L, and history depth h; each entry indicates the scaling factor α of the corresponding strategy.

Acknowledgement. The authors would like to thank the anonymous referees for helpful comments.

Copyright © 2022 by SIAM
Unauthorized reproduction of this article is prohibited

A Appendix

Proof of Theorem 3.1

Proof. For adversarial choices $d \in \mathbb{R}_{>0}$, and $o_1, o_2 \in \{-1, 1\}$, let i_0 denote the first iteration in which the agents can possibly meet if both move towards each other. That is, $i_0 := \lfloor 1 + \log_\alpha d \rfloor$ is the smallest integer such that $\alpha^{i_0} > d$.

The expected meeting time of $\mathsf{NS}(\alpha)$ decomposes into three components: the time needed for iterations $i \in \{-\infty, \ldots, i_0 - 1\}$, the time in the iteration in which rendezvous is actually achieved, and the time in between. The iterations up to iteration $i_0 - 1$ require time

$$\sum_{i=-\infty}^{i_0-1} 2\alpha^i = \frac{2\alpha^{i_0}}{\alpha - 1}.$$

In the iteration the agents achieve rendezvous, both start at their initial starting point and move straight towards each other, spending time d to meet.

As for the time in between, we know that the agents do *not* meet each other if and only if at least one agent chooses the wrong direction. Since in each iteration each agent flips a fair coin determining its direction, the probability of both agents going in the correct direction is $1/2 \cdot 1/2$, independent of the starting orientations o_1 and o_2 chosen by the adversary. Thus, the probability of not achieving rendezvous in iteration i_0 is $1 - 1/2 \cdot 1/2 = 3/4$. In this case, the probability of not achieving rendezvous in iteration $i_0 + 1$ is $3/4$ again and so on. Thus, the expected time for this component is

$$\sum_{i=i_0}^{\infty} 2\alpha^i \left(\frac{3}{4}\right)^{i+1-i_0} = \frac{6\alpha^{i_0}}{4 - 3\alpha}$$

if $\alpha < 4/3$. Thus,

$$\begin{aligned}
\rho_{\mathsf{NS}(\alpha)} &= \sup_{d \in \mathbb{R}_{>0}} \frac{\frac{2\alpha^{i_0}}{\alpha-1} + d + \frac{6\alpha^{i_0}}{4-3\alpha}}{d} = 1 + \sup_{d \in \mathbb{R}_{>0}} \left\{ \frac{\alpha^{i_0}}{d} \left(\frac{2}{\alpha - 1} + \frac{6}{4 - 3\alpha} \right) \right\} \\
&= 1 + \sup_{d \in \mathbb{R}_{>0}} \left\{ \frac{\alpha^{\lfloor 1 + \log_\alpha d \rfloor}}{d} \left(\frac{2}{\alpha - 1} + \frac{6}{4 - 3\alpha} \right) \right\} = 1 + \alpha \left(\frac{2}{\alpha - 1} + \frac{6}{4 - 3\alpha} \right) \\
&= 1 - \frac{2\alpha}{3\alpha^2 - 7\alpha + 4}
\end{aligned}$$

as claimed. ∎

Proof of Corollary 3.1

Proof. We obtain

$$\frac{\partial \rho_{\mathsf{NS}(\alpha)}}{\partial \alpha} = \frac{-2(4 - 7\alpha + 3\alpha^2) + 2\alpha(-7 + 6\alpha)}{(3\alpha^2 - 7\alpha + 4)^2} = \frac{6\alpha^2 - 8}{(3\alpha^2 - 7\alpha + 4)^2}.$$

The derivative is obviously non-positive for $\alpha \in (1, 2/\sqrt{3}]$ and non-negative for $\alpha \in [2/\sqrt{3}, 4/3)$, and the result follows. ∎

Proof of Theorem 4.1

Proof. As in the poof of Theorem 3.1, for $d \in \mathbb{R}_{>0}$ chosen by the adversary let $i_0 := \lfloor 1 + \log_\alpha d \rfloor$ be the first iteration in which the two agents can meet if they both move in the correct direction. As for the Naive Strategy, the meeting time decomposes into three components. The first component is the time for all iterations up to iteration $i_0 - 1$ which is equal to $\sum_{i=-\infty}^{i_0-1} 2\alpha^i = 2\alpha^{i_0}/(\alpha - 1)$. The third component is the time needed to achieve rendezvous in the iteration they meet, which is equal to d. We proceed to argue about the time in between, i.e, the time between iteration i_0 and the iteration when they eventually meet.

Copyright © 2022 by SIAM
Unauthorized reproduction of this article is prohibited

Each iteration $i \in \mathbb{Z}$ (if it exists) starts at time $T_i := \sum_{j=-\infty}^{i-1} 2\alpha^j$. In the time interval $[T_i, T_i + \alpha^i]$, each agent moves either α^i units to the left or α^i units to the right. Let U_i be the event that the agents have not met before iteration i and move towards each other in time interval $[T_i, T_i + \alpha^i]$. Let V_i denote the event that the agents have not met before iteration i and move in parallel in time interval $[T_i, T_i + \alpha^i]$. Finally, let W_i be the event that the agents have not met before iteration i and move away from each other in time interval $[T_i, T_i + \alpha^i]$. Furthermore, let u_i, v_i, and w_i be the probabilities that events U_i, V_i, and W_i occur, respectively.

Let $(c'_i)_{i\in\mathbb{Z}}$ and $(c''_i)_{i\in\mathbb{Z}}$ be the sequences sampled by agents 1 and 2, respectively. Using (4.2) and the fact that the agents cannot meet before iteration i_0, we obtain

$$\begin{aligned} u_{i_0} &= \mathbb{P}[c'_{i_0} o_1 = +1] \cdot \mathbb{P}[c''_{i_0} o_2 = -1] = \tfrac{1}{4}, \\ v_{i_0} &= \mathbb{P}[c'_{i_0} o_1 = -1] \cdot \mathbb{P}[c''_{i_0} o_2 = -1] + \mathbb{P}[c'_{i_0} o_1 = +1] \cdot \mathbb{P}[c''_{i_0} o_2 = +1] = \tfrac{1}{2}, \\ w_{i_0} &= \mathbb{P}[c'_{i_0} o_1 = -1] \cdot \mathbb{P}[c''_{i_0} o_2 = +1] = \tfrac{1}{4}. \end{aligned}$$

For iterations $i \in \mathbb{Z}_{>i_0}$, we describe the behavior recursively. If the event V_i occurs, the agents move in parallel. This implies that the agents will not meet in iteration i and that with probability $p^2 + (1-p)^2$ they will move in parallel again in iteration $i+1$, so that event V_{i+1} occurs. With probability $p(1-p)$, the agent that moved in the wrong direction in iteration i will change its direction while the other will stick to its direction such that event U_{i+1} occurs. With the same probability, the agent that moved in the correct direction in iteration i will change its direction while the other sticks to its direction so that event W_{i+1} occurs. With similar arguments, we obtain the following recursions:

(A.1a) $$u_{i+1} = p(1-p)v_i + p^2 w_i \qquad \text{for all } i \in \mathbb{Z}_{\geq i_0},$$

(A.1b) $$v_{i+1} = (p^2 + (1-p)^2)v_i + 2p(1-p)w_i \qquad \text{for all } i \in \mathbb{Z}_{\geq i_0},$$

(A.1c) $$w_{i+1} = p(1-p)v_i + (1-p)^2 w_i \qquad \text{for all } i \in \mathbb{Z}_{\geq i_0}.$$

We solve this system of linear recurrence relations with the method of generating functions. Let

$$G_U(z) := \sum_{i=i_0}^{\infty} u_i z^{i-i_0}, \quad G_V(z) := \sum_{i_0}^{\infty} v_i z^{i-i_0}, \quad G_W(z) := \sum_{i=i_0}^{\infty} w_i z^{i-i_0} \qquad \text{for } z \in \mathbb{R}.$$

Multiplying equation (A.1a) with z^{i+1-i_0} and summing over $i \in \mathbb{Z}_{\geq i_0}$, we obtain

$$\sum_{i=i_0}^{\infty} u_{i+1} z^{i+1-i_0} = \sum_{i=i_0}^{\infty} p(1-p) v_i z^{i+1-i_0} + \sum_{i=i_0}^{\infty} p^2 w_i z^{i+1-i_0},$$

which implies

$$G_U(z) - \tfrac{1}{4} = G_U(z) - u_{i_0} = pz(1-p)G_V(z) + p^2 z G_W(z).$$

Similarly, we obtain the equations

(A.2a) $$G_V(z) - \tfrac{1}{2} = G_V(z) - v_{i_0} = z\big(p^2 + (1-p)^2\big)G_V(z) + 2zp(1-p)G_W(z),$$

(A.2b) $$G_W(z) - \tfrac{1}{4} = G_W(z) - w_{i_0} = pz(1-p)G_V(z) + z(1-p)^2 G_W(z).$$

Rearranging (A.2b) yields

(A.3) $$G_W(z) = \frac{pz(1-p)G_V(z) + \frac{1}{4}}{1 - z(1-p)^2}$$

and plugging this equation into (A.2a), we obtain

$$G_V(z)\big(1 - zp^2 - z(1-p)^2\big) = 2zp(1-p)\frac{pz(1-p)G_V(z) + \frac{1}{4}}{1 - z(1-p)^2} + \frac{1}{2}.$$

Copyright © 2022 by SIAM
Unauthorized reproduction of this article is prohibited

Finally, we obtain

$$G_V(z)\left(1 - zp^2 - z(1-p)^2 - \frac{2p^2z^2(1-p)^2}{1-z(1-p)^2}\right) = \frac{1}{2} + \frac{zp(1-p)}{2(1-z(1-p)^2)},$$

which yields

$$G_V(z)\left(2\big(1 - zp^2 - z(1-p)^2\big)\big(1-z(1-p)^2\big) - 4p^2z^2(1-p)^2\right) = 1 - z(1-p)^2 + zp(1-p),$$

and, hence,

$$G_V(z) = \frac{1 - z(1-p)^2 + zp(1-p)}{2\big(1 - zp^2 - z(1-p)^2\big)\big(1-z(1-p)^2\big) - 4p^2z^2(1-p)^2}.$$

Plugging this equation into (A.3) then yields

$$G_W(z) = \frac{pz(1-p)}{1-z(1-p)^2} \cdot \frac{1 - z(1-p)^2 + zp(1-p)}{2\big(1 - zp^2 - z(1-p)^2\big)\big(1-z(1-p)^2\big) - 4p^2z^2(1-p)^2} + \frac{1}{4(1-z(1-p)^2}.$$

For the total expected meeting time, we obtain

$$M_{d,o_1,o_2} = \frac{2\alpha^{i_0}}{\alpha-1} + d + \sum_{i=i_0}^{\infty}(v_i + w_i)2\alpha^i = \frac{2\alpha^{i_0}}{\alpha-1} + d + 2\alpha^{i_0}\big(G_V(\alpha) + G_W(\alpha)\big).$$

This implies

$$\rho_{\mathsf{BS}(\alpha,p)} = 1 + \sup_{d\in\mathbb{R}_{>0}}\left\{\frac{\alpha^{i_0}}{d}\left(\frac{2}{\alpha-1} + 2\big(G_V(\alpha)+G_W(\alpha)\big)\right)\right\} = 1 + \frac{2\alpha}{\alpha-1} + 2\alpha\big(G_V(\alpha)+G_W(\alpha)\big),$$

since the supremum is attained for $d = \alpha^{i_0-1}$. For the given values $\alpha = 1.179$ and $p = 0.679$, one can easily check that all infinite sums used throughout the analysis do converge and all terms by which we divide are unequal zero. Moreover, the competitive ratio $\rho_{\mathsf{BS}(\alpha,p)}$ is less than 25.513.

The issue of convergence of the infinite sums in our analysis is discussed in a more general context and in more detail in the next section. ∎

Proof of Theorem 5.1

Proof. Because of the initial trembling, the agents are in the stationary distribution at iteration i_0, i.e., $\bar{\pi}_{i_0} = \bar{\pi} := \pi \otimes \pi$. From iteration i_0 on, being in a state $(\boldsymbol{s}^1, \boldsymbol{s}^2) \in S_0$ results in a rendezvous, so we have

$$\bar{\pi}_{i+1} = (\bar{\pi}_i[S_0], \bar{\pi}_i[S_+]) \cdot \begin{pmatrix} 0 \\ \bar{A}[S_+] \end{pmatrix} = \bar{\pi}_i[S_+] \cdot \bar{A}[S_+] \qquad \text{for all } i \geq i_0,$$

where $\bar{A}[S_+]$ denotes the submatrix of $\bar{A}$ formed by the rows of $\bar{A}$ indexed in S_+ and $\bar{\pi}[S]$ denotes the vector with the entries of π indexed in S. Restricting the previous equation to the coordinates in S_+ gives $\bar{\pi}_{i+1}[S_+] = \bar{\pi}_i[S_+] \cdot \bar{A}_+$, where $\bar{A}_+ = \bar{A}[S_+, S_+]$. Thus, by induction on $i \geq i_0$, the probability to be in one of the combined states $(\boldsymbol{s}^1, \boldsymbol{s}^2) \in S_+$ at iteration $i \geq i_0$ is

$$\bar{\pi}_i[S_+] \cdot \mathbf{1} \;=\; \bar{\pi}_{i_0}[S_+] \cdot \bar{A}_+^{i-i_0} \cdot \mathbf{1} \;=\; \bar{\pi}_+ \cdot \bar{A}_+^{i-i_0} \cdot \mathbf{1}.$$

Therefore, the expected time the agents spend in iterations $i \geq i_0$ not achieving rendezvous equals

$$\sum_{i=i_0}^{\infty} 2\alpha^i \cdot (\bar{\pi}_+ \cdot \bar{A}_+^{i-i_0} \cdot \mathbf{1}) \;=\; 2\alpha^{i_0}\, \bar{\pi}_+ \cdot \left(\sum_{j=0}^{\infty} \alpha^j \bar{A}_+^j\right) \cdot \mathbf{1} \;=\; 2\alpha^{i_0}\, \bar{\pi}_+ \cdot \big(I - \alpha\bar{A}_+\big)^{-1} \cdot \mathbf{1},$$

Copyright © 2022 by SIAM
Unauthorized reproduction of this article is prohibited

σ	p_σ	σ	p_σ	σ	p_σ	σ	p_σ
000100	0.638	000101	1.000	001000	1.000	001001	0.000
001010	0.980	001011	0.684	010001	0.000	010010	0.626
010100	1.000	010101	0.473	010110	1.000	010111	0.528
011010	0.578	011011	1.000	011100	1.000	011101	0.420
011110	1.000	011111	0.564	100010	0.047	100100	0.892
100101	0.733	101000	1.000	101001	0.000	101010	0.881
101011	0.805	101101	0.028	101110	0.873	101111	0.864
110010	0.000	110100	0.995	110101	0.609	110111	0.343
111001	0.000	111010	0.731	111011	1.000	111100	1.000
111101	0.282	111110	0.961	111111	0.667		

Table 3: Switching probabilities p_σ for the recurrent states $\sigma \in \mathcal{R}$ of a MARKOV STRATEGY with history depth $h = 6$, scaling factor $\alpha = 1.238$, and competitive ratio $\rho_{\mathsf{MS}(\alpha,h,p)} < 20.7533$

where the last equality requires the matrix $\alpha\bar{A}_+$ to have a spectral radius of less than 1; otherwise, the competitive ratio diverges to ∞.

As in previous sections, we obtain the expected meeting time by adding the time spent in iterations $i < i_0$, which is $\sum_{i=-\infty}^{i_0-1} 2\alpha^i = 2\frac{\alpha^{i_0}}{\alpha-1}$, and the time spent in the iteration where rendezvous is achieved, which is d. Thus, the expected meeting time can be written as

$$2\frac{\alpha^{i_0}}{\alpha-1} + d + 2\alpha^{i_0}\ \bar{\pi}_+ \cdot \left(I - \alpha\bar{A}_+\right)^{-1} \cdot \mathbf{1}.$$

Finally, dividing the above expression by d and using the fact that $\frac{\alpha^{i_0}}{d} \leq \alpha$ (by definition of i_0, we have $\alpha^{i_0-1} \leq d$) yields the result. ∎

Proof of Corollary 5.1

Proof. It is easy to verify that the above-defined Markov chain is aperiodic and irreducible. Thus, the corollary follows from Theorem 5.1. ∎

Proof of Lemma 6.1

Proof. The agents are always moving with unit speed from one peak at position $\pm d \pm \ell_{i-k}^k$ and time P_{i-k}^k to the next peak at position $\pm d \mp \ell_i^j$ and time P_i^j. Rendezvous can only be achieved while both agents are moving towards each other, i.e., when moving from a negative peak towards a positive peak, since it is impossible to rendezvous with an agent that is "running away" with the same speed it is being "chased". In all prior strategies, this situation only occurred when both agents chose the "correct" side. In the MARKOV STRATEGY WITH LOOKAHEAD, however, there is also the possibility of one agent making a lunge so long it catches the other agent on the back leg of its own lunge to the "wrong" side. (Obviously no rendezvous can occur if both agents lunge towards the wrong side.)

Let us determine the possible times M_d at which rendezvous can occur. Assume that the last peaks before rendezvous are at times $P_{t_1-k_1}^{k_1}$ and $P_{t_2-k_2}^{k_2}$, respectively, with the agents being at position $-d-\ell_{t_1-k_1}^{k_1}$ and $d+\ell_{t_2-k_2}^{k_2}$, respectively; see Figure 4. During the leg of interest, the position of both agents is thus $-d-\ell_{t_1-k_1}^{k_1}+(t-P_{t_1-k_1}^{k_1})$ and $d+\ell_{t_2-k_2}^{k_2}-(t-P_{t_2-k_2}^{k_2})$. Finding the meeting time M_d is therefore equivalent to solving the equation

$$-d-\ell_{t_1-k_1}^{k_1}+(M_d-P_{t_1-k_1}^{k_1}) \ = \ d+\ell_{t_2-k_2}^{k_2}-(M_d-P_{t_2-k_2}^{k_2}).$$

Using the identity $P_{i-k}^k+\ell_{i-k}^k = T_i$, this rewrites $2M_d = T_{t_1}+T_{t_2}+2d$. Now, set $m := \min(t_1,t_2)$ and $\delta := |t_1-t_2|$. It holds $T_{t_1}+T_{t_2} = 2T_m + 2\ell_m^\delta$. This gives:

$$M_d = T_m + \ell_m^\delta + d = P_m^\delta + d.$$

From $P_{t_1-k_1}^{k_1}$ and $P_{t_2-k_2}^{k_2}$ being the last peaks before rendezvous occurs and the fact that both agents must still be moving at each other at that time, we get $\delta = |t_1-t_2| \leq L-1$, concluding the proof. ∎

Copyright © 2022 by SIAM
Unauthorized reproduction of this article is prohibited

Proof of Lemma 6.2

Proof. We first claim that the assumption that no rendezvous occurs at some time $P_i^\delta + d$ for some $(i,\delta) < (t_1, t_2 - t_1)$, together with $d < \ell_{t_2}^{t_1+j_1-t_2}$, rules out the possibility of a rendezvous before time $P_{t_1}^{t_2-t_1} + d$. By Lemma 6.1 and our assumption, we know that rendezvous can only occur at some time $P_i^\delta + d$ with $(i,\delta) \geq (t_1, t_2 - t_1)$. To prove the claim, we shall show that $P_i^\delta \geq P_{t_1}^{t_2-t_1}$. If $i = t_1$, then the lexicographic ordering implies $\delta \geq t_2 - t_1$, and $P_i^\delta = P_{t_1}^\delta \geq P_{t_1}^{t_2-t_1}$. So let us assume that $i > t_1$. Lemma 6.1 tells us that the agents must start lunges towards each other at iterations i and $i+\delta$, respectively. But then, it must be the case that $i + \delta \geq t_1 + j_1$, as we know that one of the agents is already making a lunge spanning iterations t_1 to $t_1 + j_1 - 1$. We thus have $\delta \geq t_1 + j_1 - i$, and hence

$$P_i^\delta \;\geq\; P_i^{t_1+j_1-i} \;=\; \frac{1}{\alpha-1}\cdot(\alpha^{t_1+j_1} + \alpha^i).$$

Now, observe that $0 < d < \ell_{t_2}^{t_1+j_1-t_2}$ implies $t_1 + j_1 > t_2$. Moreover, we assumed $i > t_1$. Substituting these lower bounds for $t_1 + j_1$ and i in the previous expression yields

$$P_i^\delta \;\geq\; \frac{1}{\alpha-1}\cdot(\alpha^{t_2} + \alpha^{t_1}) \;\geq\; \frac{\alpha^{t_1}}{\alpha-1}\cdot(\alpha^{t_2-t_1} + 1) \;=\; P_{t_1}^{t_2-t_1}.$$

This proves the claim.

Let us assume (w.l.o.g.) that the agent with index 1 (i.e., such that $t_1 \leq t_2$) has starting position $-d$. Then, the agents are moving towards $-d + \ell_{t_1}^{j_1}$ and $d - \ell_{t_2}^{j_2}$ at times $P_{t_1}^{j_1}$ and $P_{t_2}^{j_2}$, respectively; see Figure 4. We have shown in the previous lemma that these trajectory legs intersect at time $t = P_{t_1}^{t_2-t_1} + d$. So, under the assumption that rendezvous has not been achieved before time t, it must occur at time t if and only if we have both $P_{t_1}^{t_2-t_1} + d < P_{t_1}^{j_1}$ and $P_{t_1}^{t_2-t_1} + d < P_{t_2}^{j_2}$, i.e., both legs complete after the candidate meeting time t. Finally, we shall see that $P_{t_1}^{t_2-t_1} + d < \min(P_{t_1}^{j_1}, P_{t_2}^{j_2}) \iff d < \min(\ell_{t_2}^{t_1+j_1-t_2}, \ell_{t_1}^{t_2+j_2-t_1})$. We first handle the inequality $P_{t_1}^{j_1} > P_{t_1}^{t_2-t_1} + d$, which rewrites

$$d \;<\; P_{t_1}^{j_1} - P_{t_1}^{t_2-t_1} \;=\; \frac{\alpha^{t_1}}{\alpha-1}(\alpha^{j_1} - \alpha^{t_2-t_1}) \;=\; \frac{\alpha^{t_2}}{\alpha-1}(\alpha^{j_1+t_1-t_2} - 1) \;=\; \ell_{t_2}^{j_1+t_1-t_2}.$$

Similarly, $P_{t_2}^{j_2} > P_{t_1}^{t_2-t_1} + d$ rewrites

$$d \;<\; P_{t_2}^{j_2} - P_{t_1}^{t_2-t_1} \;=\; \frac{\alpha^{t_2}}{\alpha-1}(\alpha^{j_2} + 1) - \frac{\alpha^{t_1}}{\alpha-1}(\alpha^{t_2-t_1} + 1) \;=\; \frac{\alpha^{t_2+j_2} - \alpha^{t_1}}{\alpha-1} \;=\; \ell_{t_1}^{t_2+j_2-t_1}.$$

Putting everything together yields the statement of the lemma. ∎

Proof of Corollary 6.1

Proof. Denote the (random) meeting time at which rendezvous occurs for a starting distance of $2\cdot d$ by M_d. Using Lemma 6.1 we can write the competitive ratio of the Markov Strategy with Lookahead in the form

$$\rho_{\mathsf{MSL}(\alpha,h,L,p)} = \frac{1}{d}\cdot \sum_{\substack{i\in\mathbb{Z}\\ 0\leq\delta\leq L-1}} \mathbb{P}[M_d = P_i^\delta + d]\cdot(P_i^\delta + d) = 1 + \frac{1}{d}\cdot \sum_{\substack{i\in\mathbb{Z}\\ 0\leq\delta\leq L-1}} \mathbb{P}[M_d = P_i^\delta + d]\cdot P_i^\delta,$$

where the last equality uses the fact that the strategy achieves rendezvous with probability 1, i.e., $\sum_{i\in\mathbb{Z},\,0\leq\delta\leq L-1} \mathbb{P}[M_d = P_i^\delta + d] = 1$. Since the P_i^δ are constants, we know that if d increases, then the competitive ratio decreases as long as the probabilities $\mathbb{P}[M_d = P_i^\delta + d]$ do not change. Reusing notation from the previous lemma, those changes happen exactly when a condition for rendezvous $d < \min(\ell_{t_2}^{t_1+j_1-t_2}, \ell_{t_1}^{t_2+j_2-t_1})$ switches from *false* to *true* due to the increasing of d, for some $t_1 \leq t_2$ and $1 \leq j_1, j_2 \leq L$. We claim that these breakpoints are all of the form ℓ_i^k for some $1 \leq k \leq L$. First, observe that $1 \leq k$ follows from $0 < \ell_i^k$. Now, we have two cases to consider: If the breakpoint is of the form $\ell_i^k = \ell_{t_2}^{t_1+j_1-t_2}$, then the superscript is $k = j_1 + t_1 - t_2 \leq j_1 \leq L$. Otherwise, the breakpoint must be of the form $\ell_i^k = \ell_{t_1}^{t_2+j_2-t_1} \leq \ell_{t_2}^{t_1+j_1-t_2}$. The

Copyright © 2022 by SIAM
Unauthorized reproduction of this article is prohibited

Algorithm 1 (TEST_RENDEZVOUS)

Input: A combined state $(\boldsymbol{s}^1, \boldsymbol{s}^2) \in \mathcal{R}^2$;
Iteration number $i \in \mathbb{Z}$ and offset $\delta \in \{0, \dots, L-1\}$;

Output: *true* if rendezvous occurs at time S_i^δ with agents 1 and 2 starting lunges in the correct direction at iterations i and $i+\delta$, respectively.

1: **if** $s^1_{-1} \vee (\neg s^1_0) \vee s^2_{\delta-1} \vee (\neg s^2_\delta)$ **then**
2:     **return** *false* ▷ Agents not starting required lunges
3: **end if**
4: $j_1 \leftarrow \max\{j \in [L] : s^1_0 = s^1_1 = \ldots = s^1_{j-1} = \texttt{True}\}$ ▷ Agent 1's lunge size
5: $j_2' \leftarrow \max\{j \in [L-\delta] : s^2_\delta = \ldots = s^2_{\delta-1+j} = \texttt{True}\}$ ▷ (lower bound) on Agent 2's lunge size
6: **return** $\left(d < \min(\ell_{i+\delta}^{j_1-\delta},\, \ell_i^{j_2'+\delta})\right)$

inequality rewrites $\alpha^{t_2+j_2} - \alpha^{t_1} \le \alpha^{t_1+j_1} - \alpha^{t_2}$, and dividing by α^{t_1} gives $\alpha^{t_2+j_2-t_1} \le 1 + \alpha^{j_1} - \alpha^{t_2-t_1} \le \alpha^{j_1} \le \alpha^L$, where we have used $\alpha^{t_2-t_1} \ge \alpha^0 = 1$ in the second inequality. The superscript is thus $k = t_2 + j_2 - t_1 \le L$, showing the claim.

We have shown that the competitive ratio is a piecewise decreasing function of d, so the worst case starting distance must be one of the breakpoints ℓ_i^k, with $1 \le k \le L$ and $i \in \mathbb{Z}$. ◻

Proof of Lemma 6.3

Proof. The proof is constructive. We test for each $\delta = 0, \dots, L-1$ whether $(\boldsymbol{s}^1, \boldsymbol{s}^2) \in S_i^\delta(d)$. If for each δ this is not the case, then we must have $(\boldsymbol{s}^1, \boldsymbol{s}^2) \in S_i^+(d)$. We claim that the routine TEST_RENDEZVOUS$((\boldsymbol{s}^1, \boldsymbol{s}^2), i, \delta)$ depicted in Algorithm 1 returns *true* if and only if rendezvous occurs at time S_i^δ with agents 1 and 2 starting lunges in the correct direction in iterations i and $i+\delta$, respectively. Then, by Lemma 6.1, we can test whether $(\boldsymbol{s}^1, \boldsymbol{s}^2) \in S_i^\delta(d)$ by running both TEST_RENDEZVOUS$((\boldsymbol{s}^1, \boldsymbol{s}^2), i, \delta)$ and TEST_RENDEZVOUS$((\boldsymbol{s}^2, \boldsymbol{s}^1), i, \delta)$; Indeed we have $(\boldsymbol{s}^1, \boldsymbol{s}^2) \in S_i^\delta(d)$ if and only if one of these tests evaluates to *true*.

To prove the claim, let $0 \le \delta \le L-1$, and assume that $(\boldsymbol{s}^1, \boldsymbol{s}^2) \notin S_i^{\delta'}(d)$ for all $\delta' < \delta$. The initial check at line 1 of the algorithm ensures that agents 1 and 2 start lunges in the correct direction at iterations i and $i+\delta$, respectively. Denote by j_1 the largest integer $j \le L$ such that $s^1_{i-1} = \texttt{True}$, $\forall i \in \{1, \dots, j\}$ and by j_2' the largest integer $j \le L - \delta$ such that $s^2_{\delta-1+i} = \texttt{True}$, $\forall i \in \{1, \dots, j\}$. Note that j_1 corresponds to our definition of the lunge size j_1 in Lemma 6.2; however, j_2' only underestimates j_2 (we would need bits for the $\delta + L$ next choices to know j_2 without ambiguity, but we only track the next L bits). Nevertheless, we claim that $(\boldsymbol{s}^1, \boldsymbol{s}^2) \in S_i^\delta(d)$ if and only if $d < \min(\ell_{i+\delta}^{j_1-\delta},\, \ell_i^{j_2'+\delta})$.

If $j_2' < L - \delta$, then j_2 coincides with the lunge size j_2 of Lemma 6.2, and the inequality $d < \min(\ell_{i+\delta}^{j_1-\delta},\, \ell_i^{j_2'+\delta})$ is simply the condition of the lemma with $i_1 = i$, $i_2 = i + \delta$, and $j_2 = j_2'$. Otherwise, we must have $j_2 \ge j_2' = L - \delta$; in that case we have

$$\ell_i^{j_2+\delta} \ge \ell_i^{j_2'+\delta} = \ell_i^L = \frac{\alpha^i}{\alpha - 1}(\alpha^L - 1) \ge \frac{\alpha^i}{\alpha-1}(\alpha^{j_1} - \alpha^\delta) = \ell_{i+\delta}^{j_1-\delta},$$

where we have used $L \ge j_1$ and $\alpha^\delta \ge 1$ in the second inequality. But then, $\min(\ell_{i+\delta}^{j_1-\delta},\, \ell_i^{j_2'+\delta}) = \min(\ell_{i+\delta}^{j_1-\delta},\, \ell_i^{j_2+\delta}) = \ell_{i+\delta}^{j_1-\delta}$, and we see that we can evaluate the condition of Lemma 6.2 even without knowing the precise value of the lunge size $j_2 \ge L - \delta$. ◻

Proof of Proposition 6.1

Proof. The first iteration at which the rendezvous can occur is i_0. We use the condition on rendezvous in Lemma 6.2 to see that i_0 is the first iteration i s.t. there exists a $1 \le \delta \le L$ with $\ell_i^\delta > d$. Since $\ell_i^{\delta'} > \ell_i^{\delta''}$ for $\delta' > \delta''$, we can assume $\delta = L$. Furthermore, $\ell_i^\delta = \alpha^i \cdot \ell_0^\delta$. Thus, i_0 is the smallest i s.t. $\alpha^i \ell_0^L > d$, so $i_0 = \lfloor \log_\alpha {}^{d}/_{\ell_0^L} \rfloor + 1$.

Copyright © 2022 by SIAM
Unauthorized reproduction of this article is prohibited

i_1 is the first iteration such that the classifications $S_i^*(d)$ are the same for all $i \geq i_1$. Using Lemma 6.2 for the classification, we see that i_1 is the first iteration i such that $\ell_i^\delta > d \iff \ell_{i'}^\delta > d$ for all $i' \geq i$ and $1 \leq \delta \leq L$. We can choose i' to be arbitrarily large, thus we can always make the left hand side true. Also $\ell_i^{\delta'} > \ell_i^{\delta''}$ for $\delta' > \delta''$. Thus, i_1 is the first iteration i such that $\ell_i^0 > d$. This means $i_1 = \lfloor \log_\alpha d \rfloor + 1$. ☐

Proof of Theorem 6.1

Proof. As a technical matter, we first prove that the probability of achieving rendezvous is 1 for $d \in \{\ell_0^1, \cdots, \ell_0^L\}$. Assume this is not the case and add an artificial absorbing state `a` to the Markov chain, into which all agents achieving rendezvous are routed. By our assumption, this Markov chain is not absorbing, so there must be a recurring state `r`, i.e., a state from which the (single!) artificial absorbing state is not accessible. So, the class of communicating states of `r` does not contain `a` and consists of only recurring states. It thus has a stationary distribution not including `a`. Therefore, after omitting the entry for `a`, the state vector of that stationary distribution is an eigenvector of $\bar{A}_{i_1}(d)$ with eigenvalue 1. In particular, the spectral radius of the matrix $\alpha\bar{A}_{i_1}(d)$ is at least $\alpha > 1$, yielding a contradiction to the theorem's prerequisite.

Assume d is fixed. Because of the initial trembling, the agents must be in the stationary distribution at iteration i_0, i.e., $\bar{\pi}_{i_0} = \bar{\pi} := \pi \otimes \pi$. Then, by construction, the probability distribution over the combined states at iteration $i \geq i_0$ is $\bar{\pi}_i = \bar{\pi}_{i_0} \cdot \bar{A}_i(d)^{i-i_0}$, where $\bar{\pi}_{i_0} := \pi \otimes \pi$.

As in the proof of Corollary 6.1, we write the competitive ratio

$$\rho_{\mathsf{MSL}(\alpha,h,L,p)} = \frac{1}{d} \cdot \sum_{\substack{i \in \mathbb{Z} \\ 0 \leq \delta \leq L-1}} \mathbb{P}[M_d = P_i^\delta + d] \cdot (P_i^\delta + d) = 1 + \frac{1}{d} \cdot \sum_{\substack{i \in \mathbb{Z} \\ 0 \leq \delta \leq L-1}} \mathbb{P}[M_d = P_i^\delta + d] \cdot P_i^\delta, \tag{A.4}$$

where the last equality uses the the fact that rendezvous is achieved with probability 1.

By definition of i_0 we have $P[M_d = P_i^\delta + d] = 0$ for all iterations $i < i_0$. Then, for $i_0 \leq i < i_1$ we have $\mathbb{P}[M_d = P_i^\delta + d] = \bar{\pi}_i \cdot \boldsymbol{\chi}_{S_i^\delta(d)}$ and recall $P_i^\delta = \alpha^i \frac{\alpha^\delta + 1}{\alpha - 1}$. Thus,

$$\sum_{\substack{i_0 \leq i < i_1 \\ 0 \leq \delta \leq L-1}} \mathbb{P}[M_d = P_i^\delta + d] \cdot P_i^\delta = \alpha^{i_0} \sum_{i=0}^{i_1 - i_0 - 1} \alpha^i \sum_{\delta=0}^{L-1} \frac{\alpha^\delta + 1}{\alpha - 1} \cdot \bar{\pi}_i \cdot \boldsymbol{\chi}_{S_i^\delta(d)}. \tag{A.5}$$

For the iterations $i \geq i_1$ we proceed analogously to the analysis of Theorem 5.1 and get

$$\begin{aligned}
\sum_{\substack{i_1 \leq i \\ 0 \leq \delta \leq L-1}} \mathbb{P}[M_d = P_i^\delta + d] \cdot P_i^\delta &= \sum_{j=0}^{\infty} \alpha^{i_1 + j} \sum_{\delta=0}^{L-1} \frac{\alpha^\delta + 1}{\alpha - 1} \cdot \bar{\pi}_i \cdot \boldsymbol{\chi}_{S_i^\delta(d)} \\
&= \alpha^{i_1} \sum_{j=0}^{\infty} \alpha^j \sum_{\delta=0}^{L-1} \frac{\alpha^\delta + 1}{\alpha - 1} \cdot \bar{\pi}_{i_1} \cdot \bar{A}_{i_1}(d)^j \cdot \boldsymbol{\chi}_{S_{i_1}^\delta(d)} \\
&= \alpha^{i_1} \cdot \bar{\pi}_{i_1} \cdot (I - \alpha \bar{A}_{i_1}(d))^{-1} \cdot \left(\sum_{\delta=0}^{L-1} \frac{\alpha^\delta + 1}{\alpha - 1} \cdot \boldsymbol{\chi}_{S_{i_1}^\delta(d)} \right),
\end{aligned} \tag{A.6}$$

where we used the fact that $\alpha\bar{A}_{i_1}(d)$ has spectral radius at most 1 in the last equation.

Next, we prove that the competitive ratios for $d = l_p^\delta$ and $d = l_q^\delta$ coincide for $p, q \in \mathbb{Z}$ and $\delta \in \{0, \cdots, L-1\}$. According to Equation (A.4), the competitive ratio is given by one plus the sum of the terms (A.5) and (A.6) divided by d. The denominator d changes by a factor of α^{q-p} when it gets changed from l_p^δ to l_q^δ. Due to Proposition 6.1, the leading factors α^{i_0} and α^{i_1} in (A.5) and (A.6), respectively, also change by a factor of α^{q-p} when d gets changed from from l_p^δ to l_q^δ. It remains to show that the rest of (A.5) and (A.6) are not affected by changing d from l_p^δ to l_q^δ. Due to Proposition 6.1, i_0 and i_1 shift by the same amount and the stationary distribution is independent from the starting distance. That $\chi_{S_i^\delta(l_p^\delta)} = \chi_{S_{i-p+q}^\delta(l_q^\delta)}$ for $i \in \mathbb{Z}$ follows from the classification of states in Lemma 6.2 and the definition of the l_i^δ.

The Theorem follows with the restriction on possible worst case starting distances of Corollary 6.1. ☐

Copyright © 2022 by SIAM
Unauthorized reproduction of this article is prohibited

σ	p_σ	σ	p_σ	σ	p_σ	σ	p_σ
000001	0.803	000010	0.000	000011	1.000	000100	0.239
000101	0.566	000110	1.000	000111	0.000	001000	0.405
001001	1.000	001010	0.112	001011	0.362	001100	1.000
001101	0.000	001110	0.000	001111	0.352	010000	0.835
010001	0.343	010010	0.000	010011	0.525	010100	1.000
010101	0.000	010110	0.675	010111	1.000	011001	0.000
011010	0.000	011011	0.549	011100	0.000	011101	0.992
011110	0.438	011111	0.796	100000	1.000	100001	0.240
100010	0.676	100011	0.295	100100	0.144	100110	0.000
100111	1.000	101000	0.379	101001	0.000	101010	0.258
101011	0.262	101100	1.000	101101	1.000	101110	1.000
101111	0.489	110000	0.792	110001	0.325	110010	0.000
110100	0.000	110101	0.282	110110	0.033	110111	0.634
111000	0.356	111001	0.000	111010	0.686	111011	0.732
111100	1.000	111101	0.250	111110	0.252	111111	0.404

Table 4: Switching probabilities p_σ for the recurrent states $\sigma \in \mathcal{R}$ of a Markov Strategy with Lookahead with history depth $h = 6$, maximal lunge size $L = 6$, scaling factor $\alpha = 1.18684$, and competitive ratio $\rho_{\mathsf{MSL}(\alpha,h,L,p)} < 13.9262$

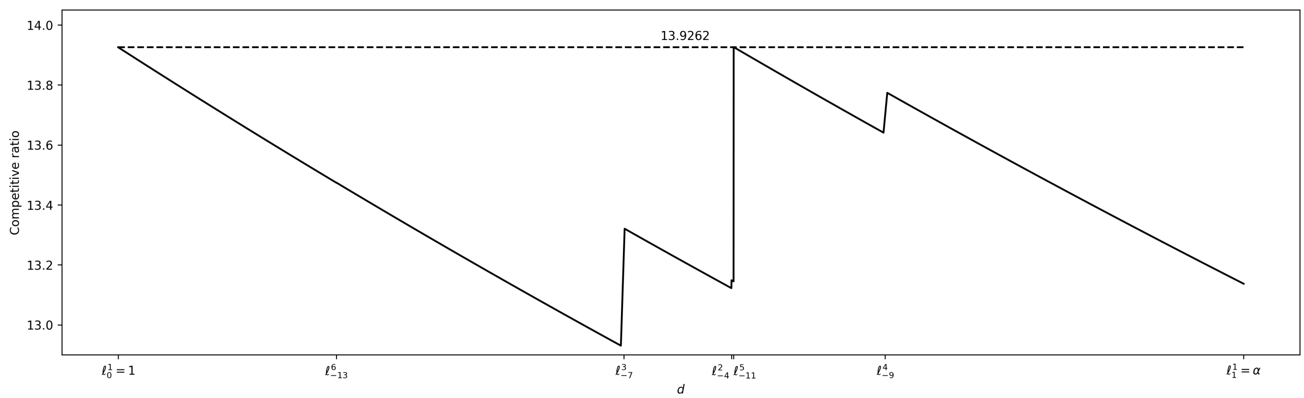

Figure 5: Competitive ratio as a function of the unknown starting distance (solid line) for $d \in [1, \alpha)$, and overall competitive ratio $\rho_{\mathsf{MSL}} < 13.9262$ of the Markov Strategy with lookahead and lunges of size $L = 6$ of Table 4 (dashed line); as expected from Corollary 6.1, the breakpoints occur at values of the form $d = \ell_i^k$ for $k \in \{1, \ldots, 6\}$.

References

[1] Steve Alpern. The rendezvous search problem. *SIAM Journal on Control and Optimimization*, 33:673–683, 1995.

[2] Steve Alpern. Rendezvous search: A personal perspective. *Operations Research*, 50:772–795, 2002.

[3] Steve Alpern and Anatole Beck. Asymmetric rendezvous on the line is a double linear search problem. *Mathematics of Operations Research*, 24:604–618, 1999.

[4] Steve Alpern and Anatole Beck. Pure strategy asymmetric rendezvous on the line with an unknown initial distance. *Operations Research*, 48(3):498–501, 2000.

[5] Steve Alpern and Shmuel Gal. Rendezvous search on the line with distinguishable players. *SIAM Journal on Control and Optimization*, 33:1270–1276, 1995.

[6] Steve Alpern and Shmuel Gal. Searching for an agent who may or may not want to be found. *Operations Research*, 50:311–323, 2002.

[7] Steve Alpern and Shmuel Gal. *The theory of search games and rendezvous*, volume 55. Springer Science & Business Media, 2006.

[8] Edward J. Anderson and Skander Essegaier. Rendezvous search on the line with indistinguishable players. 33(6):1637–1642, 1995.

[9] Ricardo A. Baezayates, Joseph C. Culberson, and Gregory J.E. Rawlins. Searching in the plane. *Information and computation*, 106(2):234–252, 1993.

[10] Vic Baston. Note: Two rendezvous search problems on the line. *Naval Research Logistics*, 46:335–340, 1999.

Copyright © 2022 by SIAM
Unauthorized reproduction of this article is prohibited

[11] Vic Baston and Shmuel Gal. Rendezvous on the line when the players' initial distance is given by an unknown probability distribution. *SIAM Journal on Control and Optimization*, 36(6):1880–1889, 1998.

[12] Qiaoming Han, Donglei Du, Juan Vera, and Luis F. Zuluaga. Improved bounds for the symmetric rendezvous value on the line. *Operations Research*, 56(3):772–782, 2008.

[13] Ming-Yang Kao, John H Reif, and Stephen R Tate. Searching in an unknown environment: An optimal randomized algorithm for the cow-path problem. *Information and Computation*, 131(1):63–79, 1996.

[14] Deniz Ozsoyeller, Andrew Beveridge, and Volkan Isler. Symmetric rendezvous search on the line with an unknown initial distance. *IEEE Transactions on Robotics*, 29(6):1366–1379, 2013.

Copyright © 2022 by SIAM
Unauthorized reproduction of this article is prohibited