RENDEZVOUS SEARCH ON LABELLED NETWORKS

STEVE ALPERN

Abstract. Two players are independently placed on a commonly labelled network X. They cannot see each other but wish to meet in least expected time. We consider continuous and discrete versions, in which they may move at unit speed or between adjacent distinct nodes, respectively. There are two versions of the problem (asymmetric or symmetric) depending on whether or not we allow the players to use different strategies. After obtaining some optimality conditions for general networks, we specialize to the interval and circle networks. In the first setting, we extend the work of J. V. Howard; in the second we prove a conjecture concerning the optimal symmetric strategy.

1. Introduction

The rendezvous search problem [1] asks how two unit-speed players can optimally meet after they are randomly placed in a known region X according to known independent distributions. We consider here the case where the search region X is a network, that is, a finite union of arcs which are joined only at their ends. For networks, we will say that the players meet when their positions coincide exactly. The notion of an optimal meeting will be one occurring in least expected time, or more generally with least expected cost, where cost is increasing in time. We assume of course that the players cannot see each other prior to meeting.

In rendezvous problems it is usually assumed that after the players' initial placement in X they have some uncertainty about their position or about the direction they are facing. In [1], this uncertainty is formalized in terms of a given group of symmetries of X. However in this paper such ambiguities are avoided by taking the simpler assumption that the network X is fully labelled, and that this labelling is common to both players. For those readers familiar with the more general formulation of rendezvous search, this assumption is equivalent to saying that the given symmetry group is the trivial one consisting of only the identity isomorphism. The assumption that the network is fully labelled has the consequence that if player I is initially placed at a point $x \in X$, he simply chooses a path $f_x(t)$ in X with $f_x(0) = x$ and $|f_x(t_2) - f_x(t_1)| \leq |t_2 - t_1|$. A strategy for a player is simply a collection of such paths, one for each possible starting point x. If players I and II choose respective strategies f and g and their starting points turn out to be x and y, then they will meet at time

$$T_{x,y}(f,g) = \min\{t : f_x(t) = g_y(t)\}.$$

Consequently if their initial distributions are μ and ν then their expected meeting time will be given by

$$\tilde{T}(f,g) = \int_{X\times X} T_{x,y}(f,g)\, d_\mu(x)\, d\nu(y),$$

assuming that the strategies f and g have sufficient regularity (say piecewise continuous on X) so that $T_{x,y}$ is integrable. If the players can get together before the initial placement, then they can decide together on the strategy pair (f,g) so as to minimize $\tilde{T}(f,g)$. This version of the problem is called the *asymmetric rendezvous problem* and denoted $\Gamma^a(X,\mu,\nu)$. The minimal expected meeting time is called the *asymmetric rendezvous value*, which we formally define as

$$R^a(X,\mu,\nu)=\min_{f,g}\tilde{T}(f,g).$$

Another version of the problem is concerned with the case where rendezvous strategies are to be printed in a guidebook that say all hikers carry with them. When they become separated and wish to meet, they look up their search region in the guidebook and follow the strategy that is recommended. So we consider what strategy to recommend, given that both players will follow it. We call this problem the *symmetric rendezvous problem* $\Gamma^s(X,\mu)$, which we consider only in the case of a common initial distribution $\mu=\nu$. The corresponding symmetric rendezvous value R^s is defined by the optimization problem

$$R^s(X,\mu)=\min_{f}\tilde{T}(f,f).$$

Following their introduction in [1], both versions of the rendezvous problem have been extensively studied (see [4],[5], [6], [8], [9], [13], [14]), but mainly with non-trivial symmetry groups so that the players cannot choose their exact paths in X, but only equivalence classes of paths. The only paper to address the fully labelled version of rendezvous search is that of J. V. Howard [12], where the asymmetric problem is solved for the circle (with a common initial uniform distribution) and both versions are solved (with a common uniform or increasing distribution) for the interval. Howard [12] showed that for the circle the optimal asymmetric solution is for the players to move at unit speed in opposite directions until they meet (as suggested in [1]). For the interval, we define the Howard strategy as the one which when starting at a point $x\in[0,1]$ goes to the right until time $(1-x)/2$ and then goes to the left, always at full speed 1. This paper extends [12] by showing that the Howard strategy is uniquely optimal for any strictly increasing distributions p,q and any cost function which is increasing and convex in meeting time T. The techniques used here combine extensions of Howard's techniques with some new general ones.

The approach taken by Howard in his pioneering article [12] was to approximate the interval (enlarged to $[0,2(n-1)]$) by an interval graph with node set $\{0,1,\ldots,2(n-1)\}$ and a discrete time set $t=0,1,2,\ldots$ on which players move to adjacent nodes. While the study of rendezvous on graphs is not new (see [3] and the classic original article of Anderson and Weber [7]), Howard [12] ensured that his discretization method was a good model for the continuous version. In the continuous version, players who transpose their positions along an interval must have met (Intermediate Value Theorem), but if the graph version ends when players are at the same node they may transpose between adjacent nodes without meeting. To avoid this possibility, Howard requires that the initial distribution p (he considered only one, to be used for both players) was concentrated on the even numbered nodes. Since the graph distance between such nodes is always even, this ensures that at even times t both players are at even nodes and at odd times they are both at odd nodes, and so they cannot pass each other without meeting. He then solved the discrete asymmetric rendezvous problem on the interval graph $I[n,p]$ and showed that as n goes to infinity his solution converges to an optimal solution to the continuous problem for the interval.

We extend Howard's discretization to an arbitrary network in the following manner. Along each arc in the continuous network X, we add an odd number of additional nodes of degree two. This ensures that the graph length of each of the original arcs is even and consequently that all circuits of the resulting discrete graph (based on the original nodes and additional nodes) have even length. We call a node 'even' if it is an original node of X or a new node which is an even distance from an original node. We further require that the distributions over the new nodes which describe the initial placement of the players are concentrated on the even nodes. We will call such graphs, which satisfy both conditions, *H-networks*, since they model the meeting properties of paths on continuous networks. This paper is mainly concerned with rendezvous search on H-networks and consequent results for continuous networks.

The paper is organized as follows. In Section 2 we define the discrete rendezvous problems (asymmetric and symmetric) on an H-network, and show that optimal strategies for either problem must move along geodesics of the network between potential meetings with the other players. Such strategies are called geodesic. In Section 3 we specialize to the symmetric problem and prove the existence of optimal symmetric strategies with the property (called 'sticky') that if $f_x(t) = f_y(t)$ then for $t' > t$ we also have $f_x(t') = f_y(t')$ (the paths f_x and f_y must stick together after they meet). In Section 4 we consider an example of asymmetric rendezvous on a particular H-network where no sticky strategy pair is optimal. In Section 5 we extend the techniques and results of Howard [12] for the interval H-network. We show that a strategy suggested by Howard is uniquely optimal for strictly increasing initial distributions, even when the cost function is not simply time but rather any increasing convex function of time. For the symmetric problem with an initial equiprobable distribution, we show that *any* symmetric strategy which is sticky and geodesic is optimal. In Section 6 we consider an H-network model of the circle and show that there is always an optimal symmetric strategy for which some (odd) node of the network is avoided by both players, regardless of their initial placement. Consequently this node can be removed without increasing the symmetric rendezvous value. However its removal leaves an interval H-network, for which the solution is known. This shows that the symmetric rendezvous value of the continuous circle of unit circumference has symmetric rendezvous value 1/3, a result conjectured in [1].

2. Rendezvous on Labeled H-Networks

Let X denote a graph with nodes $\mathcal{N}$. Assume that all circuits have even length, so that (by Theorem 2.4 of [11]) we can partition the nodes into $\mathcal{N} = \mathcal{N}_e \cup \mathcal{N}_o$ such that every path in X alternates between $\mathcal{N}_e$ and $\mathcal{N}_o$. Such a (bipartite) graph, together with two given probability distributions p and q on the even node set $\mathcal{N}_e$, will be called an *H-network*. An example of such a network is given below in Figure 1, which will be used later to illustrate a counterexample. The even nodes $\mathcal{N}_e = \{A, B, C, D, E, F, G\}$ are labelled and indicated by small disks, while the odd nodes are unlabelled and indicated only by short lines.

Figure 1: An H-Network with $\mathcal{N}_e = \{A, B, C, D, E, F, G\}$

Given an H-network (X, p, q), the asymmetric rendezvous problem $\Gamma^a(X, p, q)$ is defined as follows. At time $t = 0$, two players I and II are placed independently on the even nodes $\mathcal{N}_e$, according to the distributions p, q. At each time $t = 1, 2, \ldots,$ they move to a distinct adjacent node. This model insures that they cannot pass each other without meeting (transpose their positions between adjacent nodes) because the nodes they occupy will always have the same parity. A strategy s for a player specifies for each even node $i \in \mathcal{N}_e$ a path $s_i(t)$ in the network X which starts at node i. This means that $s_i(t)$ and $s_i(t+1)$ are distinct adjacent nodes of the network, and $s_i(0) = i$. The interpretation is that a player using strategy s who is initially placed at node i will follow the path $s_i(t)$. Let $\mathcal{S}$ denote the set of all strategies. If the players adopt the strategies f and g, and they are initially placed at nodes $i, j \in \mathcal{N}_e$, then they will meet at time

$$T_{i,j}(f, g) = \min\{t : f_i(t) = g_j(t)\}. \tag{2.1a}$$

Note that $T_{i,i} = 0$. We will only consider strategy pairs (f, g) for which all the $T_{i,j}$ are finite, and for these we will define $t_{f,g} = \max T_{i,j}(f, g)$. The expected meeting time for two players adopting the strategy pair (f, g) is

$$\tilde{T}(f, g) = \sum_{i,j \in \mathcal{N}_e} T_{i,j}(f, g)\, p_i q_j. \tag{2.2}$$

Any strategy pair minimizing (2.2) will be called *optimal.* The minimum expected time will be called the asymmetric rendezvous value of the H-network (X, p, q) and denoted $R^a (X, p, q)$.

The symmetric rendezvous problem $\Gamma^s (X, p)$ is the same except that we assume that $q = p$ (players are placed with a common distribution) and that $g = f$ (the players must adopt a common strategy). A strategy f minimizing $\tilde{T} (f, f)$ will be called an optimal symmetric strategy, and the corresponding value of $\tilde{T}$ will be called the symmetric rendezvous value of (X, p) and denoted $R^s (X, p)$. Obviously we have that $R^a (X, p, p) \leq R^s (X, p)$. Note that for labelled spaces we do not need to consider mixed strategies, even for the symmetric problem. The symmetric problem corresponds to the situation where the players cannot be distinguished and must read a recommended strategy from a common book on rendezvous theory. For the symmetric problem we shall adopt the same notations as defined above for the asymmetric case, except that there will be a single strategy and a single distribution. For example t_f denotes the last time $t_{f,f}$ that players using f may meet.

To illustrate the notion of a strategy we consider the interval graph based on nodes $0, 1, 2, \ldots, 6$. Figure 2 shows two strategies, drawn with locations on the horizontal axis and time increasing upwards on the vertical axis. Note that search paths are defined only for even location. In both cases the end agents (starting at 0 and 6) go towards the center, reaching it at time 3. For the strategy on the left we use a dotted line to indicate the path of the agent starting at location 4 to distinguish it from the strategy starting at location 2. (The strategy on the right can be drawn in a single style because it is 'sticky' in a sense defined precisely in the following section).

Figure 2: Two strategies for the interval H-network

The meeting time matrices for the symmetric form of these strategies, and of their pairing, are given below. (The rows and columns correspond to the even nodes

0,2,4,6 in the usual order.)

$$T(h,h) = \begin{pmatrix} 0 & 3 & 2 & 3 \\ 3 & 0 & 1 & 2 \\ 2 & 1 & 0 & 3 \\ 3 & 2 & 3 & 0 \end{pmatrix}, T(g,g) = \begin{pmatrix} 0 & 1 & 3 & 3 \\ 1 & 0 & 3 & 3 \\ 3 & 3 & 0 & 2 \\ 3 & 3 & 2 & 0 \end{pmatrix}, T(h,g) = \begin{pmatrix} 0 & 1 & 3 & 3 \\ 1 & 0 & 3 & 3 \\ 3 & 3 & 0 & 2 \\ 3 & 3 & 2 & 0 \end{pmatrix}$$

$$\tilde{T}(h,h) = \frac{28}{16}, \qquad \tilde{T}(h,h) = \frac{30}{16}, \qquad \tilde{T}(h,h) = \frac{30}{16}.$$

With reference to a given strategy s under consideration, it will be useful to think of the paths $s_i(t)$ as paths of a collection of *agents* of a player, with one agent starting at each even node. The formula (2.2) for $\tilde{T}$ can be thought of as the expected time for an agent of the player I to meet an agent of player II. It will often be useful to consider the set of g-agents (agents of player II) that a particular agent $i \in \mathcal{N}_e$ of f (of player I) has met by time t, and this subset of $\mathcal{N}_e$ will be denoted by

$$\mathcal{M}_i(f,g,t) = \{m : T_{i,m}(f,g) \leq t\}.$$

Lemma 1. *Let (f,g) an optimal strategy pair for $\Gamma^a(X,p,q)$, let $i \in \mathcal{N}_e$ be any agent, and let t_0 denote any integer time. Then the sub-path $f_i(t)$, $t \geq t_0$, minimizes the expression*

$$\sum_{k \in \sim\mathcal{M}_i(s,t)} q_k T_{i,k}(f,g), \tag{2.3}$$

representing the (un-normalized) expected time to meet the remaining agents of the other player. The same result holds for the player II strategy g.

Proof. Suppose $\bar{f}$ is equal to f except for coordinate i and times $t \geq t_0$. Then

$$\tilde{T}(f,g) - \tilde{T}(\bar{f},g) = \sum_{k \in \sim\mathcal{M}_i(s,t)} q_k T_{i,k}(f,g) - \sum_{k \in \sim\mathcal{M}_i(s,t)} q_k T_{i,k}(\bar{f},g).$$

So if the conclusion of the lemma is false, then the right hand side of this expression will be positive, and consequently (by the left side) the assumed optimality of f will be contradicted. ∎

The above lemma can be used to give a necessary condition for optimality of an agent f_i with respect to its full strategy pair (f,g). Let $\mathcal{T}_i$ denote the set of times when agent f_i meets agents of g (of player II). That is, $\mathcal{T}_i = \cup_j \{T_{i,j}(f,g)\}$. Suppose the t_0 and t_1 are consecutive times in $\mathcal{T}_i$. We claim that if (f,g) is optimal then agent f_i must move along a geodesic of the network X from to t_0 to time t_1. We formalize this in the following definition.

Definition 1. *A pair (f,g) is called a geodesic strategy if for any even node $i \in \mathcal{N}_e$, and any consecutive times t_0 and t_1 in the meeting time set $\mathcal{T}_i$ of agent f_i, we have*

$$d_{\mathcal{N}}(f_i(t_0), f_i(t_1)) = t_1 - t_0, \tag{2.4}$$

where $d_{\mathcal{N}}$ denotes the graph distance (length of shortest path) in the network X. The corresponding condition must also hold for g.

The concept of a geodesic strategy is illustrated by the two strategies drawn for the interval network in Figure 2. The symmetric strategy (h,h) drawn on the left is geodesic but (g,g) is not, roughly because of the turn taken by the agent of g starting at 4 at time $t = 1$, between $t_0 = 0$ and $t_1 = 2$. This strategy can be improved (and made geodesic) by having this agent move to location 5 at time $t=1$, and then staying with the agent starting at 6. This technique is used in the following result, which will be useful in limiting the set of strategies that we must consider as candidates for optimality. See Theorem 5 for a surprising converse to this result for interval networks.

Theorem 1. *A necessary condition for a strategy pair (f,g) to be optimal for the asymmetric problem is that it is geodesic. A necessary condition for a strategy f to be optimal for the symmetric problem is that (f,f) is geodesic.*

Proof. We prove the asymmetric claim first. Suppose that (f,g) is an optimal strategy pair which fails the geodesic condition (2.4) for f_i and times t_0 and t_1. Let k satisfy $T_{i,k}=t_1$. In other words $f_i(t_1)=g_k(t_1)$, or agent f_i meets agent g_k (as well as possibly other agents of player II) at time t_1. Observe that the integers $d_{\mathcal{N}}(f_i(t_0), f_i(t_1))$ and t_1-t_0 must have the same parity, since they are both the lengths of paths (which move on every step) from $f_i(t_0)$ to $f_i(t_1)$. Consequently we have that

$$d_{\mathcal{N}}(f_i(t_0), f_i(t_1)) \leq t_1 - t_0 - 2. \tag{2.5}$$

We now modify the pair (f,g) only for f_i for times t strictly between t_0 and t_1, and call the modified pair (f',g'). By (2.5) be may ensure that $f_i'(t_1-2)=g_k(t_1)$ and consequently that $f'(t_1-1)=g_k(t_1-1)$. Hence $T_{i,k}' \leq t_1-1 = T_{i,k}-1$. Since none of the meeting times have increased, the time (2.3) has decreased, and consequently (f,g) cannot have been optimal. The same proof works for the symmetric case. ∎

3. Sticky strategies for symmetric rendezvous

We now restrict our attention to the symmetric rendezvous problem $\Gamma^s(X,p)$ and show that it is sufficient in this case to consider sticky strategies in which agents who meet subsequently stay together.

The symmetric problem $\Gamma^s(X,p)$ may be thought of as a game among $\#(\mathcal{N}_e)$ agents rather than between two players. In the game interpretation, we may consider that each agent i in $\mathcal{N}_e$ (who starts at even node i) wishes to minimize the expected time to meet the other agents. In other words, his utility can be described by $U_i=\sum_{k\neq i} p_k T_{ik}$. The symmetric form of Lemma 1 says that any optimal symmetric strategy for the rendezvous problem $\Gamma^s(X,p)$ is a subgame perfect Nash equilibrium for the game version. Since we will need this form later on, we state it below.

Lemma 2. *Let f be an optimal symmetric strategy for $\Gamma^s(X,p)$, let $i\in\mathcal{N}_e$ be any agent, and let t_0 denote any integer time. Then the sub-path $f_i(t)$, $t\geq t_0$, minimizes the expression*

$$\sum_{k\in\sim\mathcal{M}_i(s,t)} p_k T_{i,k}(f),$$

representing the (un-normalized) expected time to meet the remaining agents.

We formalize the notion of a sticky strategy as follows.

Definition 2. *We will say that a strategy $s \in \mathcal{S}$ is sticky at time t, if for any $i, j \in \mathcal{N}_e$,*

$$s_i(t) = s_j(t) \ \text{ and } t' > t \Rightarrow s_i(t') = s_j(t'). \tag{3.1}$$

If a strategy is sticky at all times t, we will simply call it sticky.

The strategy drawn on the left (called h) in Figure 2 is not sticky, because the agents starting at locations 2 and 4 meet and then separate. The strategy drawn on the right is sticky, and this is why it can be drawn in a single style without ambiguity. The strategy h on the left can be modified to make it sticky by choosing one of the two continuations from location 3 at time 1 for *both* players to follow. This idea is used Lemma 4 below.

If a strategy is sticky then agents at a common location will have met exactly the same set of agents (though perhaps at different times). In particular, we have the following.

Lemma 3. *Suppose that the symmetric strategy f is sticky at times $t = 1, 2, \ldots, t_0 - 1$. Then for any $i, j \in \mathcal{N}_e$, $f_i(t_0) = f_j(t_0)$ implies that $\mathcal{M}_i(f, t_o) = \mathcal{M}_j(f, t_o)$.*

Proof. Let k belong to $\mathcal{M}_i(f, t_o)$. By the symmetry of the statement it is enough to prove that k also belongs to $\mathcal{M}_j(f, t_o)$. First suppose that $T_{i,k}(f) < t_0$. Then by the sticky assumption for time $T_{i,k}(f)$, we know that $f_k(t_0) = f_i(t_0)$. But $f_i(t_0) = f_j(t_0)$ and so $k \in \mathcal{M}_j(f, t_o)$. On the other hand if $T_{i,k}(f) = t_0$ then it follows from the definition of $T_{i,k}$ that $f_i(t_0) = f_j(t_0)$, so the conclusion follows in this case as well. ∎

Suppose that at some time two agents of a player, say f_i and f_j, are at the same location and have met the same set of agents of the other player. Then clearly these two agents face exactly the same optimization problem for the future, so any continuation strategy optimal for one is optimal for the other. In particular, we have

Lemma 4. *Suppose that f is an optimal symmetric strategy for $\Gamma^s(X, p)$ and that for some integer time t_o its agents $i, j \in \mathcal{N}_e$ satisfy:*

1. *$f_i(t_0) = f_j(t_0)$, and*
2. *$\mathcal{M}_i(f, t_o) = \mathcal{M}_j(f, t_o)$.*

Then if the strategy f' is equal to f except for agent f_j and time $t \geq t_0$, where $f'_j(t) \equiv f_i(t)$, the strategy f' is also optimal for $\Gamma^s(X, p)$.

Proof. According to Lemma 1, both f_i and f_j, $t \geq t_0$, must be optimal for the subgame determined by the location $f_i(t_0)$ and the past meetings $\mathcal{M}_i(f, t_o)$. Consequently changing $f_j(t)$ to $f_i(t)$, $t \geq t_0$, will not change the expected meeting time. ∎

We can now prove the main result of this section, that there are always optimal sticky strategies for symmetric rendezvous on an H-network.

Theorem 2. *For any symmetric rendezvous problem $\Gamma^s(X, p)$ on an H-network (X, p), there exists an optimal symmetric strategy which is sticky.*

Proof. We prove this by induction on t_0 for the statement that there exists an optimal strategy which is sticky for times up to t_0. We first show that any optimal strategy f can be modified to produce an optimal strategy which is sticky at time $t=1$. Suppose $f_i(1)=f_j(1)=x$ for distinct $i,j\in\mathcal{N}_e$. Then $\mathcal{M}_i(f,1)=\mathcal{M}_j(f,1)=\{k:f_k(1)=x\}$. Consequently we may apply the previous lemma with $t_0=1$ to change f to an optimal strategy which in which $f_i(t)=f_j(t)$ for $t\geq 1$. Continuing in this manner produces an optimal strategy which is sticky for time 1. Now suppose that f is sticky for times $t=1,2,\ldots,t_0-1$. We will modify it so that it is sticky at time t_0, completing the induction proof. So suppose that for some new agents which we again call i and j, we have $f_i(t_0)=f_j(t_0)$. It follows from Lemma 2 that $\mathcal{M}_i(f,t_o)=\mathcal{M}_j(f,t_o)$, and the required modification is now possible according to Lemma 3. ∎

4. An Example

In this section we analyze the particular H-network shown earlier in Figure 1, which illustrates many of the general ideas in the previous section. The asymmetric rendezvous problem on this network also provides a counterexample to an extension of Theorem 2, on the optimality of sticky strategies, to the asymmetric setting.

We will analyze the asymmetric problem on this network, though we will consider only a symmetric initial distributions $p=q$. In particular, define p^ε by $p^\varepsilon(A)=p^\varepsilon(B)=p^\varepsilon(C)=\frac{1}{3}-\frac{(\varepsilon+\varepsilon^2)}{3}$, $p^\varepsilon(D)=\varepsilon$, and $p^\varepsilon(E)=\varepsilon^2$. That is, most of the probability is equally on A,B, and C, a small amount is on D, and a much smaller amount is on E. The probability of starting at any other node is zero. The analysis given in this section shows that, up to labelling the two players, the unique optimal asmmetric strategy pair for small ε is as follows.

$$
\begin{array}{ll}
f_A=(A,\_,B,\_C,\_,G) & g_A=(A,\_,C,\_G)\\
f_B=(B,\_,C,\_G) & g_B=(B,\_,A,\_,C,\_,G)\\
f_C=(C,\_,A,\_,B,\_C) & g_C=(C,\_,B,\_,A,\_,C)\\
f_D=(D,\_,C,\_,A,\_C) & g_D=(D,\_C,\_,B,\_C)\\
f_E=(E,\_,F,\_,G,\_,C) & g_E=(E,\_,F,\_,G,\_,C)
\end{array}
\tag{4.1}
$$

In particular, note that $f_B(2)=f_D(2)=C$, but $f_B(4)=G\neq A=f_D(4)$. So the unique optimal symmetric strategy is not sticky. This shows that Theorem 2 cannot be extended to the asymmetric setting.

We will analyze this network by first considering the two simpler subnetworks determined by the even arc sets $\{A,B,C\}\subset\{A,B,C,D\}\subset\mathcal{N}_e$.or equivalently, initial distributions which are concentrated on these subsets.

First consider the distribution $\left(\frac{1}{3},\frac{1}{3},\frac{1}{3}\right)$ on $\{A,B,C\}$. For intuition, recall that Howard [12] showed that for the asymmetric problem on the continuous circle the optimal strategy is for one player to go clockwise and the other counterclockwise. The discrete version of this fact is easy to see. Observe that of the nine entries in the meeting time matrix for $\{A,B,C\}$ there will be three 0's (corresponding to common initial placements of the players) and at most three 1's. Consequently no strategy pair can be better than one with three 0's, three 1's and three 2's in its meeting time matrix. Such a matrix is achieved uniquely by the discrete version of the clockwise-counterclockwise strategy, namely:

$$
\begin{array}{|c|c|}\hline
f_A=(A,\_,B) & g_A=(A,\_,C)\\\hline
f_B=(B,\_,C) & g_B=(B,\_,A)\\\hline
f_C=(C,\_,A) & g_C=(C,\_,B)\\\hline
\end{array},\ \text{with } T=\begin{pmatrix}0&1&2\\2&0&1\\1&2&0\end{pmatrix},\ \tilde{T}=\frac{9}{9}=1.
\tag{4.2}
$$

Here the underscore represents the unlabelled intervening odd node, and the ordering on $\mathcal{N}_e$ (or the rows and columns of T) is alphabetical. Note that as there are only finitely many strategy pairs, any other pair (except for reversal of players I and II) has $\tilde{T} \geq 1+\delta$, for some fixed $\delta > 0$. Next put a mass ε on the node D which is small with respect to δ. (This gives total mass $1+\varepsilon$, but we could normalize if we wanted.) Our variational approach ensures that the strategy beginnings given in (4.2) should not change. Thus after f_D starts with $(D, \_, C)$, it has met g_A, and it can best meet g_B and g_C by the continuations listed in the table below. For full strategies labelling of the odd nodes is unnecessary, but to avoid ambiguity for partial strategies (defined up to certain times) we label odd nodes by putting a '$+$' on the node clockwise of a given node.

$f_A = (A, \_, B, B+)$	$g_A = (A, \_, C,)$
$f_B = (B, \_, C,)$	$g_B = (B, \_, A, C+)$
$f_C = (C, \_, A, \_, B)$	$g_C = (C, \_, B, \_, A)$
$f_D = (D, \_, C, \_A)$	$g_D = (D, \_C, \_, B)$

, with $T = \begin{pmatrix} 0 & 1 & 2 & 3 \\ 2 & 0 & 1 & 2 \\ 1 & 2 & 0 & 4 \\ 2 & 3 & 4 & 0 \end{pmatrix}$.

Note that f_B and f_D meet at node C at time $t = 2$. At this time f_B has met all the agents of the other player (g_A, g_B, g_C, g_D) while f_D has not yet met g_C, who he will meet at node A at time $t = 4$. Observe that f_B and g_A are defined only up to time $t = 2$ since the second (B) row and first (A) column of T have a maximum of 2.

If we now place a mass γ at E which is very small with respect to ε, the above strategies will still represent the optimal way of starting. Clearly the optimal strategy for either agent (of player I or II) at E is to go towards C ($f_E = g_E = (E, \_, F, \_, G, \_C)$) and for the agents starting at A, B, C, D it is optimal (after they have met each other) to go to C and then towards E, starting as soon as the other agents have been met. (The node E was placed sufficiently far from C to make this obvious.) This gives the unique optimal continuations of the above strategies as stated above in (4.1).

5. Rendezvous on an Interval

In this section we extend the work of Howard [12] on the interval H-network $I[n,p,q]$ designed by Howard to model the continuous interval $[0,1]$. The node set is given by $\mathcal{N} = \{0, 1, \ldots, 2(n-1)\}$, with $\mathcal{N}_e = \{0, 2, \ldots, 2(n-1)\}$ and $\mathcal{N}_o = \{1, 3, \ldots, 2n-3\}$. The distributions p and q are concentrated on $\mathcal{N}_e$. Nodes i and j are adjacent if $|i-j| = 1$. To simplify the notation, we will denote by f_i the path of the strategy f that begins at the even node $2i$, so that $f_i(0) = 2i$. (This differs from our general notation.) For this problem we will also consider a generalization of the basic cost function T by allowing an increasing function $C(T)$. (The basic problem has $C(t) = t$.)

As observed by Howard, we may assume that an agent starting at either end (0 or $2(n-1)$) of the interval will proceed in the direction of the center $n-1$ until he reaches it at time $n-1$. That is, we restrict the strategy set $\mathcal{S}$ so that for all $s \in \mathcal{S}$ we have $s_0(t) = t$ and $s_{n-1}(t) = 2(n-1) - t$. We will call these two possible agents the 'left sweeper' and the 'right sweeper' respectively, as in [12]. It is also easy to see that any agent who meets either sweeper should subsequently follow the sweeper path, so we make this restriction as well.

A strategy is determined by the 'steps' $\triangle f_k(t) \equiv f_k(t+1) - f_k(t) = \pm 1$. This is interpreted as saying that an agent starting at location $2k$ will move to the right (or left) between times t and $t+1$ if $\triangle f_k(t) = +1$ (or -1).

We picture below in Figure 4 three strategies for the case $n = 8$ with lines drawn for the paths of each of the 8 agents $g_0, \ldots, g_7$. For each of these the search space is the drawn horizontally and time moves upwards. For each even starting point a line is drawn which in each upward step slants either to the left or right. The left and right sweepers form with the base an right triangle. When two lines meet they stay together, as all these strategies are sticky. Also note that the lines only turn when they meet another line, as all these strategies are geodesic.

Figure 3: Strategies for the interval network

The top strategy, suggested by Howard [12] and denoted by $\hat{s}$, is to move to the right each period until meeting the right sweeper, and then to move left with this sweeper. For each starting point k, this strategy is given in our step notation by

$$\triangle \hat{s}_k = (+1, +1, \dots, +1, -1, -1, \dots, -1),$$

where there are $n-1-k$ 'plus ones' required to meet the right sweeper, followed by k 'minus ones'. The middle strategy corresponds to the simple statement, 'go to the center node $n-1$'. But as staying still is not allowed, this is followed by an oscillation between $n-1$ and the node n. The bottom strategy is what we call the 'greedy strategy', as agents maximize the number of meetings at time $t=1$, and then those at time $t=3$, and so on. This family is well defined when n is a power of 2 ($n = 2^3$ in the figure). For the smaller parameter $n=4$, the meeting time matrices $T_{i,j},\ i,j = 0,1,2,3$, are given below for the symmetric version of the Howard, central, and greedy, strategies. Note that in all cases the meeting times $T_{i,j}$ sum to 28 and so the expected meeting time is 28/16. (Theorem 5 below will give an for an explanation of this common value, which is also minimal.) Note that the first two strategies have the same set of meeting times, and so they will give the same expected value for any cost $C(T_{i,j})$, while the greedy strategy will be better if the players face a large cost if they fail to meet at time $t=1$.

$$\begin{pmatrix} 0 & 3 & 3 & 3 \\ 3 & 0 & 2 & 2 \\ 3 & 2 & 0 & 1 \\ 3 & 2 & 1 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 3 & 3 & 3 \\ 3 & 0 & 1 & 2 \\ 3 & 1 & 0 & 2 \\ 3 & 2 & 2 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 1 & 3 & 3 \\ 1 & 0 & 3 & 3 \\ 3 & 3 & 0 & 1 \\ 3 & 3 & 1 & 0 \end{pmatrix}$$

Howard showed that the symmetric pair $(\hat{s}, \hat{s})$ is optimal in the asymmetric problem for any n and any equal nondecreasing distributions p and q. Using a variant of his original argument, we establish the following result on *unique* optimality for p and q which need no longer be equal but must be strictly increasing. We also show that unique optimality persists for more general cost functions.

Theorem 3. *For any n, the Howard strategy pair $(\hat{s}, \hat{s})$ is uniquely optimal for the asymmetric interval rendezvous problem $I[n,p,q]$, as long as both p and q are strictly increasing. This results holds more generally for any cost function C of the meeting time, as long as C is nondecreasing and convex.*

Proof. Let (f,g) be any strategy pair other than $(\hat{s}, \hat{s})$. We must show that (f,g) is not optimal. We consider first the simpler case $C(t) = t$. Let k be the smallest integer for which it is *not* true that $f_k = g_k = \hat{s}_k$. This will be true for some k with $1 \le k \le n-2$. By renumbering the agents (interchanging f and g), if necessary, we may assume that

$$f_i = \hat{s}_i,\ i = 0, \dots, k-1, \text{ and} \tag{5.1}$$

$$g_k \neq \hat{s}_k. \tag{5.2}$$

By (5.2), there is a largest time r such that $\Delta g_k(r) = -1$ while $\Delta \hat{s}_k(r) = +1$. Note that r is the last time that agent g_k (who starts at $2k$ and uses strategy g) moves left before meeting the right sweeper. For example if $n=9$, $k=3$, and (omitting the 1's)

$$\begin{aligned} \Delta g_3 &= (+,+,-,+,-^*,+,+,-), \\ \Delta \hat{s}_3 &= (+,+,+,+,+\ \ ,-,-,-), \end{aligned} \tag{5.3}$$

then we would have $r = 4$ (remember we begin with time coordinate zero), with the defining minus indicated by a *. Define a strategy g' which equals g except for coordinate $g_k(t)$, $t > r$. Define $\Delta g'_k(r) = +1$ and $\Delta g'_k(t) = \Delta g_k(t)$ for $t > r$ (until g'_k reaches the right sweeper when we must set $\Delta g'_k(t) = -1$). Thus in the example of equation (5.3) we have

$$\Delta g'_3 = (+, +, -, +, +^*, +, -, -) .$$

See Figure 4, where g_3 and g'_3 are drawn in solid lines. Note that the region between these two paths is a rectangle. The base of the rectangle is always a single slanted line, while the height consists of m slanted lines, where m will be determined later. The reason why g' is better than g against f is that while it meets the agent of f who moves along the left side of the rectangle m time units later than g does, it meets each of the m rightmost agents of f at least one time unit earlier. Since agents to the right are assumed to be more probable than those on the left, this results in a shorter expected meeting time with a random agent of f. Consequently the pair (f, g) cannot be optimal, completing the proof for $C(t) = t$. (In the example illustrated in Figure 4, although agent g'_3 meets the agent starting at location $2 = 2(L-1)$ exactly $m = 2$ time units later than agent g_3 does, he meets each of the two $(= m)$ agents starting at location 14 and 16 at least one time unit earlier than g_3 does.)

Figure 4: Plots of g_k and g'_k .

We now make the above argument more rigorous, as well as incorporating a non-decreasing convex function C of meeting time. Observe that from time r until the time $r + m$ when g_k reaches the right sweeper, g'_k will be one unit to the right of g_k. At all other times the paths determined by g'_k and g_k will coincide. The determination of m, the height of the rectangle between the two paths, depends on how far to the left g_k has gone by time r (the number of -1's preceding $\triangle g_k(r)$).

For some integer L, $1 \le L \le k$, agent g_k at time r will be at the same location as agent f_L (or $\hat{s}_L$). Since $g_k(r) = -1$, g_k will meet f_{L-1} in the next time period $r+1$. However g'_k will not meet f_{L-1} until f_{L-1} meets the right sweeper at time $n-1-L$. So the meeting with g'_k is $m \equiv (n-1-L)-(r+1)$ time units later than with g_k, namely at time $r+1+m$ rather than at time $r+1$. Consequently

$$C\left(T_{L-1,k}(f,g)\right) - C\left(T_{L-1,k}(f,g')\right) = C(r+1) - C(r+1+m) \tag{5.4}$$

Next observe that when any of the rightmost m agents f_{n-j}, $j = 1, \ldots, m$, meet agent g'_k, the agent g_k will be one unit to the left of g'_k and so cannot meet f_{n-j} for at least one more time unit. Furthermore, the meeting of agent f_{n-j} with g'_k will occur at some time $r+m+1-j+i$, where i is a non-negative integer, since $r+m+1-j$ is the earliest possible meeting. Consequently for each $j = 1, \ldots, m$, we have

(5.5)
$$C\left(T_{n-j,k}(f,g)\right) - C\left(T_{n-j,k}(f,g')\right) \ge C(r+m+2-j+i) - C(r+m+1-j+i).$$

Since we are assuming that C is convex, the right hand side of the above inequality is nondecreasing in i, and consequently we have

(5.6)
$$C\left(T_{n-j,k}(f,g)\right) - C\left(T_{n-j,k}(f,g')\right) \ge C(r+m+2-j) - C(r+m+1-j), \text{ for } j = 1, \ldots, m.$$

The meeting point pairs representing the actual difference in cost (5.5) and the estimated difference in cost (5.6) are joined by the thick lines in Figure 5.

For all other values of i, $T_{i,k}(f,g) - T_{i,k}(f,g') = 0$. Also, for all $j \ne k$ we have $T_{i,j}(f,g) - T_{i,j}(f,g') = 0$. Combining (5.4) and (5.6) with the expected meeting time formula (2.2), with $T_{i,j}(f,g)$ replaced by the more general cost $C\left(T_{i,j}(f,g)\right)$, we obtain

$$\begin{aligned}
&\tilde{T}(f,g,p,q) - \tilde{T}(f,g',p,q) \\
&\ge q_k \left[p_{L-1}\left(C(r+1) - C(r+1+m)\right) + \sum_{j=1}^{m} p_{n-j}\left(C(r+m+2-j) - C(r+m+1-j)\right) \right] \\
&> q_k \left[p_{L-1}\left(C(r+1) - C(r+1+m)\right) + p_{n-m} \sum_{j=1}^{m} \left(C(r+m+2-j) - C(r+m+1-j)\right) \right] \\
&\quad \text{(since } p \text{ is strictly increasing and } C \text{ is nondecreasing)} \\
&= q_k \left[p_{L-1}\left(C(r+1) - C(r+1+m)\right) + p_{n-m}\left(C(r+m+1) - C(r+1)\right) \right] \\
&= q_k \left(p_{n-m} - p_{L-1}\right)\left(C(r+m+1) - C(r+1)\right) \\
&> 0, \text{ since } p \text{ is strictly increasing and } C \text{ is nondecreasing.}
\end{aligned}$$

Consequently the strategy pair (f,g) cannot be optimal with respect to C, p, and q. ∎

Figure 5: Meeting times for g_k and g'_k.

The above argument is a modification of that given in [12], in which g_k was immediately improved to $\hat{s}_k$. Howard's argument gives a generally larger improvement, but the exact value of the improvement is harder to evaluate explicitly.

The following result of Howard (for $p = q$) shows that $(\hat{s}, \hat{s})$ is always optimal (though no longer uniquely so - see Theorem 5) even if the initial distributions are not *strictly* increasing. We prove it here as a consequence of the uniqueness result Theorem 3.

Corollary 1. *For any n, the Howard strategy pair $(\hat{s}, \hat{s})$ is optimal for the interval asymmetric rendezvous problem $I[n, p, q]$, as long as both p and q are nondecreasing.*

Proof. Suppose on the contrary that for some nondecreasing distributions p, q, and some strategy pair (f, g) we have

$$\tilde{T}(\hat{s}, \hat{s}, p, q) - \tilde{T}(f, g, p, q) > 0.$$

The formula (2.2) shows that the expected meeting time $\tilde{T}$ is continuous in p and q for fixed strategies. Consequently for all p' and q' in sufficiently small neighborhoods of p and q in the $(n-1)$ −simplex of probability distributions, we also have that

$$\tilde{T}(\hat{s}, \hat{s}, p', q') - \tilde{T}(f, g, p', q') > 0.$$

However this contradicts the previous theorem, since such neighborhoods contain strictly increasing distributions. ∎

Howard [12] used Corollary 1 to obtain the following solution to the continuous version of the problem.

Theorem 4. *Let $\Gamma^a([0,1], p(x))$ denote the asymmetric continuous rendezvous problem on the unit interval $[0,1]$, in which each player is independently initially placed according to a common nondecreasing density function p. An optimal strategy pair is the following continuous version of $\hat{s}$. Each agent moves at unit speed to the right until meeting the right sweeper, and moves with him to the left at unit speed. The associated rendezvous value is*

$$R^a([0,1], p, p) = \frac{1}{2}\left(1 - \int_0^1 (1 - F_p(x))^2 \, dx\right),$$

where F_p is the distribution function corresponding to the density p.

Proof. We briefly outline Howard's proof. For a more rigorous proof see Howard [12] Suppose there was a strategy pair for $\Gamma^a([0,1], p(x))$ giving a lower expected rendezvous time than the continuous Howard strategy pair of moving right and then following the right sweeper. The interval $[0,1]$ can be scaled to $[0, 2(n-1)]$ and continuous strategies can be approximated by strategies that move as in the discrete problem. These discrete strategies in $I[n, \tilde{p}, \tilde{p}]$ would do better than $(\hat{s}, \hat{s})$, where $\tilde{p}$ is a discrete version of the density p. However this is impossible by Corollary 1. ∎

It is interesting to note that the optimal strategy pair for this asymmetric problem consists of two identical strategies. Consequently these strategies must also be optimal for the symmetric problem. This is an example where the symmetric rendezvous value equals the asymmetric rendezvous value.

Now let us specialize to the case of the uniform density $\bar{p}(x) = 1$ on $[0,1]$, where we have

$$R^a([0,1]) = \frac{1}{2}\left(1 - \int_0^1 (1-x)^2 \, dx\right) = \frac{1}{3}$$

Another strategy which is optimal in this case (in the symmetric or asymmetric version) is for each player to go directly to the middle point $1/2$ and wait there. The optimality is simply a matter of checking that the expected time is $1/3$, which is just the expected value of the maximum of two independent random variables which are uniformly distributed on $[0, 1/2]$.

For the equiprobable discrete probability distribution $\bar{p}_i \equiv 1/n$, the symmetric versions of the three strategy types shown in Figure 3 all have the same expected meeting time. For the Howard and the central strategies, this can be calculated as $\frac{2}{n^2}\sum_{k=1}^{n-1} k^2$ on the original discrete interval $\{0, \ldots, 2(n-1)\}$ or

$$\frac{1}{n^3}\sum_{k=1}^{n-1} k^2 = \sum_{k=0}^{n-1}\left(\frac{1}{n}\right)\left(\frac{k}{2n}\right)^2 \to \int_0^1 x^2 \quad dx = \frac{1}{3}$$

for a normalized interval H-network with nodes $\left\{0, \frac{1}{2n}, \frac{2}{2n}, \ldots, 1\right\}$ and corresponding time periods $\frac{1}{2n}$. The fact that all these symmetric strategies have the same (minimal) expected meeting time is a consequence of their being sticky and geodesic, as seen by the following result. (Howard [12] also mentions a large number of strategies which are optimal for the equiprobable distribution.) We remind the reader that we are only considering strategies in which the sweeper who start at the ends of the interval always go towards the center until they meet.

Theorem 5. *A necessary and sufficient condition that a sticky strategy be optimal for the symmetric rendezvous problem on the interval network* $I[n, \bar{p}]$ *, where* $\bar{p}$ *is the equiprobable distribution* $\bar{p}_i \equiv 1/n$*, is that it is geodesic. In particular (sufficiency) every sticky geodesic strategy is optimal for the symmetric problem.*

Proof. This result can probably be proved by direct analysis, but it is easier for us to simply observe that the proof is in fact a special case of the argument already given in the proof of Theorem 3, taking C to be the identity, $p = q = \bar{p}$, and $f = g$ (which we call g). The proof of Theorem 3 modified a given strategy pair (f, g) to (f', g') with $\tilde{T}(f, g', p, q) < \tilde{T}(f, g, p, q)$ for strictly increasing p and q. We will show that in the symmetric case of that proof, the same modification of g to g' has the property $\tilde{T}(g', \bar{p}) = \tilde{T}(g, \bar{p})$. Since after a finite number of these modifications g has become the optimal Howard strategy $\hat{s}$, we know that the original g must have been optimal. To establish that $\tilde{T}(g', \bar{p}) = \tilde{T}(g, \bar{p})$, it suffices to show that (5.6) becomes an equality. That is, we must show that each of the last m agents meets g'_k exactly one time unit before he meets g_k. This is the same as showing that none of the last m agents g_j, $j = n - m, \ldots, n - 1$, can turn to the right (a -1 followed by a $+1$ in $\triangle g_j$) when it meets the path g'_k (which is *not* an agent of g).

We now prove that such a turn cannot take place, assuming that g is sticky and geodesic. So as in the proof of Theorem 3 suppose that g_k is the first agent not following the Howard strategy $\hat{s}$, and that g_k makes its first turn to the left at a point $A = (2k + r, r)$, which means that this occurs after it has made r steps to the right. (Since g is sticky we will have $L = k - 1$, and consequently $m \equiv (n - 1) - (k - 1) - (r + 1) = n - 1 - k - r$. This situation is drawn in Figure 5, with g_k drawn with a thick line and g'_k drawn with a dotted line. We will show that the pictured turn CDE taken by agent g_j is not possible.

Figure 6: Idea from proof of Theorem 5

First observe that since g_k is part of a geodesic strategy and it turns at A, another agent (of g) must move from B to A. Similar reasoning extends the line AB back to the bottom of the picture at $2(k+r)$, so that we know that $g_{k+r}(r) = g_k(r) = 2k+r$. Since we are assuming that g is sticky, no agent of g can go to the right from A. Suppose now that g_j is the leftmost of the m rightmost agents mentioned above which makes a turn (in this case CDE) when it reaches g'_k (which is not an agent of g). The same reasoning as before based on the assumption that g is geodesic shows that there must be an agent of g going from A to D, and in fact that g_k itself must go through D. But we have already shown that this is impossible. This argument shows that when such an agent g_j reaches g'_k he must continue going to the left to meet g_k in the next time period. ∎

6. Symmetric rendezvous on a labelled circle

We now consider the symmetric rendezvous problem on the circle H-network $Circ[n,p]$ with node set $\mathcal{N}^C = \{0,1,\ldots,2n-1\}$ and even nodes $\mathcal{N}_e^C = \{0,2,\ldots,\ 2(n-1)\}$. Nodes i and j are adjacent if $|i-j| = 1 \pmod{2n}$. We will call the direction from i to $i+1$ clockwise. This problem is analyzed in [3] without the H-network structure (initial placement at any node) and without a common notion of direction. The assumptions made in this article make the problem much easier for the players, as we shall see. Also note that under our assumptions the players can rendezvous at least as well on the circle as the interval, as they could choose not to traverse some agreed location, thus 'cutting' the circle. The main idea of this section is the demonstration that in some sense they cannot do *better* on the circle than on the interval. An example of a symmetric sticky strategy for $Circ[6,p]$ is drawn in Figure 7. The location $2n-1 = 2(6)-1 = 11$ is also drawn as location -1, with an identified dotted line above each. The agents starting at locations 0 and 2 move from the left side to the right side of the picture at respective times 1 and 3.

Figure 7: A symmetric strategy on $Circ[6, p]$.

It is easy to see from this picture that a counter-clockwise rotation (leftwards in the picture) of 4 units results in a picture that would represent a valid strategy (without identifications) for the interval H-network $I\,[6]$. This is pictured in Figure

8.

Figure 8: The same strategy with relabelling.

The following result shows that the relabelling that created Figure 8 from Figure 7 is always attainable.

Theorem 6. *For every symmetric sticky strategy on the circle H-network $Circ[n]$, there is an even circular relabelling of the locations so that the resulting strategy is feasible for the interval network $I[n]$.*

Proof. Let f be any sticky strategy for $\Gamma^s\left(Circ[n]\right)$. (We remove the initial distribution p from our notations here, as it is irrelevant for this argument.) The proof gives an explicit algorithm for constructing the relabelling. First note that for each time $t = 0, 1, \ldots, t_f = \max T_{i,j}\left(f\right)$ there is a partition $P\left(t\right)$ of $\mathcal{N}_e$ into sets of type $\mathcal{M}_i\left(f, t\right)$. Specifically, two agents i and $j \in \mathcal{N}_e$ will be in the same element of the partition $P\left(t\right)$ if $T_{i,j}\left(f\right) \leq t$. Since f is assumed to be sticky, the partition $P\left(t-1\right)$ refines that of $P\left(t\right)$, with $P\left(0\right)$ consisting of all singletons and $P\left(t_f\right)$ having the single element $\mathcal{N}_e$. In the case of the circle, each element of $P\left(t\right)$ is a connected set (an arc). If $P\left(t-1\right) = A_1 \cup A_2 \cup \cdots \cup A_N$ with the arcs A_i written in say clockwise order around the circle, then in $P\left(t\right)$ the arc A_i cannot be joined with *both* A_{i-1} and A_{i+1}. Consequently the penultimate partition $P\left(t_f - 1\right)$ must consist of exactly *two* nonempty arcs A and B. (For the strategy pictured in Figure 7, we have $A = \{4, 6, 8, 10, 0\}$ and $B = \{2\}$.) By renaming these arcs if necessary, we may assume that $f_i\left(t_f - 1\right) = k$ for all $i \in A$, $f_j\left(t_f - 1\right) = k + 2 \bmod \left(2n\right)$ for all $i \in B$, and that $f_i\left(t_f\right) = k + 1$ for all $i \in \mathcal{N}_e$. (In Figure 7, the value of k is 8.) In this case we renumber the nodes so that $A = \{0, \ldots, 2m\}$ and $B = \{2m+2, \ldots, 2\left(n-1\right)\}$ for some m. ∎

It is worth noting that in the algorithm given in the above proof, the trajectories of the resulting strategy on the interval may not look exactly like that given in

Figure 8 in that the locations covered may not stay within the original nodes. For example if the strategy f begins by having all agents go clockwise for six time units, then clearly in any relabelling the trajectory starting at $2(n-1)$ (location 10 in Figure 8) will go to the right at least as far as $2(n-1)+6$. Of course if f is optimal this cannot happen, and in any case we will allow such suboptimal strategies for the interval.

We now examine the consequences of Theorem 6 in the presence of an initial distribution p. For such a distribution on even nodes of the interval, $\mathcal{N}_e = \{0, 2, \ldots, 2(n-1)\}$, we define the k'th rotation of p as p^k, where $p^{(k)}(i) = p(i+k)$, with arithmetic modulo $2n$. So for example in Figure 8 the probability of a start at location 0, $p^4(0)$, should be equal to the probability of starting on the circle at location 4. We can now apply Theorem 6 to obtain the following form of the result that rendezvous on the circle is not easier than rendezvous on the interval.

Corollary 2. *For any distribution p on $\{0, 2, \ldots, 2(n-1)\}$, there is an integer $k \in \{0, 2, \ldots, 2(n-1)\}$ such that*

$$R^s(Circ[n,p]) \geq R^s\left(I\left[n,p^k\right]\right). \tag{6.1}$$

Proof. By Theorem 2 there is an optimal strategy f for the symmetric rendezvous problem on $R^s(Circ[n,p])$ which is sticky. By the previous result we can relabel the nodes of the circle by an even rotation k so that so that the resulting strategy is feasible for the interval with the initial distribution p^k. The required inequality now follows from the assumed optimality of f. ∎

If we now apply the fact we observed in the introduction to this section, that players on the circle could agree to avoid the location $2k+1$, and hence play the game on $I\left[n,p^k\right]$, we obtain the following complete reduction of symmetric rendezvous on the circle to that on the interval.

Corollary 3. $R^s(Circ[n,p]) = \min_k R^s\left(I\left[n,p^k\right]\right)$.

Proof. Playing optimally after agreeing to avoid $2k+1$ results in an expected rendezvous time of $R^s\left(I\left[n,p^k\right]\right)$, so for each k we have $R^s(Circ[n,p]) \leq R^s\left(I\left[n,p^k\right]\right)$. Consequently the required equality follows from (6.1). ∎

Observe that the equiprobable (uniform) distribution $\bar{p}$ is invariant under the k-rotation, that is, $\bar{p}^k = \bar{p}$. Consequently for all k we have $R^s\left(I\left[n,\bar{p}^k\right]\right) = R^s(I[n,\bar{p}])$, yielding the following.

Theorem 7. *For the equiprobable distribution $\bar{p}$ on $\{0, 2, \ldots, 2(n-1)\}$, we have*

$$R^s(Circ[n,\bar{p}]) = R^s(I[n,\bar{p}]).$$

So in particular, an optimal symmetric strategy for $\Gamma^s(Circ[n,\bar{p}])$ is to go to some agreed location and then oscillate. (Other optimal strategies can be obtained from Theorem 5.)

Using the same approach as in the proof of Theorem 4 (or in the more detailed version given by Howard in [12]), we obtain the following continuous analog establishes the optimality of the strategy for symmetric rendezvous on a labelled circle conjectured in [1].

Theorem 8. *Suppose two players are placed independently and uniformly on the unit circumference circle and allowed to move at unit speed. They cannot see each* *other but have a common labelling of all the points on the circle. An optimal symmetric rendezvous strategy is to move at unit speed, via the shorter of the two possible arcs, to a commonly agreed location. Another optimal strategy is to move clockwise at unit speed until reaching a location at time t which is a distance t counter-clockwise of an agreed point, and then to move at unit speed counterclockwise. There are many other optimal strategies. The symmetric rendezvous value equals 1/3.*

We note again that the optimal solution to the *asymmetric* version of this problem, as demonstrated by Howard [12], is for one player to go clockwise and the other to go counter-clockwise.

References

[1] S. Alpern. The rendezvous search problem. *SIAM J. Control & Optimization*, **33** (1995), 673-683.

[2] S. Alpern. Asymmetric rendezvous search on the circle. *Dynamics and Control*, **10**, No. 1, in press..

[3] S. Alpern, V. Baston, and S. Essegaier. Rendezvous search on a graph. *Journal of Applied Probability*, **36** (1999), No. 1, 223-231.

[4] S. Alpern and Anatole Beck. Asymmetric rendezvous on the line is a double linear search problem. *Mathematics of Operations Research*, **24** (1999), No. 3, 604-618.

[5] E. J. Anderson and S. Essegaier (1995). Rendezvous search on the line with indistinguishable players, *SIAM Journal of Control and Optimization*, **33**, 1637-1642.

[6] E. J. Anderson and S. Fekete (1998). Two-dimensional rendezvous search. *Proceedings, 14'th ACM Symposium on Computational Geometry* (1998), 365-373 (extended abstract). *Operations Research*, to appear (1999).

[7] E. J. Anderson and R. R. Weber. The rendezvous problem on discrete locations. *Journal of Applied Probability* **28**, (1990) 839-851.

[8] V. J. Baston. Two rendezvous search problems on the line, *Naval Research Logistics* **46** (1999) 335-340.

[9] V. J. Baston, S. Gal . Rendezvous on the line when the players' initial distance is given by an unknown probability distribution. *SIAM Journal of Control and Optimization* **36** (1998), no. 6, 1880-1889.

[10] S. Gal. Rendezvous search on the line. *Operations Research,* **47** (1999), 974-976.

[11] F. Harary. Graph Theory. Adison-Wesley, London (1972).

[12] J. V. Howard. Rendezvous search on the interval and circle. *Operations Research,* **47**, No. 4 (1999), 550-558.

[13] W. S. Lim. A rendezvous-evasion game on discrete locations with joint randomization. *Advances in Applied Probability* **29** (1997), 1004-1017.

[14] L. C. Thomas and P. B. Hulme. Searching for targets who want to be found. *Journal of the O. R. Society* **48** (1997), Issue 1, Pages 44 - 50.

Department of Mathematics, The London School of Economics, Houghton Street, London WC2A 2AE, United Kingdom

E-mail address: `s.alpern@lse.ac.uk`