

THE RENDEZVOUS PROBLEM ON DISCRETE LOCATIONS

E. J. ANDERSON AND
R. R. WEBER,* *University of Cambridge*

Abstract

Two friends have become separated in a building or shopping mall and wish to meet as quickly as possible. There are n possible locations where they might meet. However, the locations are identical and there has been no prior agreement where to meet or how to search. Hence they must use identical strategies and must treat all locations in a symmetrical fashion. Suppose their search proceeds in discrete time. Since they wish to avoid the possibility of never meeting, they will wish to use some randomizing strategy. If each person searches one of the n locations at random at each step, then rendezvous will require n steps on average. It is possible to do better than this: although the optimal strategy is difficult to characterize for general n , there is a strategy with an expected time until rendezvous of less than $0.829 n$ for large enough n . For $n = 2$ and 3 the optimal strategy can be established and on average 2 and $8/3$ steps are required respectively. There are many tantalizing variations on this problem, which we discuss with some conjectures.

DYNAMIC PROGRAMMING; SEARCH PROBLEMS

1. The rendezvous problem

In this paper we introduce a new class of problems concerned with two players who wish to meet as quickly as possible. The problem we describe has some unusual features and gives rise to a bewildering number of related problems, many of which appear to be interesting in their own right. In the last part of this paper we give a brief review of some of these related problems, together with a number of conjectures and open problems.

Perhaps the best way to introduce the problem is to suppose that two people become separated in a building, a department store for example, and wish to meet again. If they have no prearranged plan to deal with this situation, it will not be easy to decide what to do. Should one of them stay where he or she is until found by the other? And, if so, for how long should they wait (since perhaps they have both decided to stay still)? The simplest mathematical formulation which captures some of the features of this situation is to suppose that there are n different locations in the building, and

Received 18 September 1989; revision received 9 January 1990.

* Postal address for both authors: University of Cambridge, Department of Engineering, Management Studies Group, Mill Lane, Cambridge CB2 1RX, UK.

that the two players will find each other if they are in the same location, but not otherwise. Since the two players have not agreed a strategy between them, we can think of the locations as identical and not labelled in any way. If the two players are intelligent they may well decide to play a mixed strategy which, if played by both players, gives the shortest expected time until they meet. To keep things simple, we suppose that movement between locations is instantaneous, that the players begin their searches at the same time, that time is split into periods and that each player visits only one location in each period.

This is a type of search problem with two players. However, it is unusual to have such a problem in which the players are co-operating to find each other. More typically search problems with two players either involve one player moving randomly (see for example Weber (1986)) or they involve a pursuer and an evader, who seek to minimize and maximize respectively the time until they meet (see for example Gal (1980), Alpern and Asic (1986)). A problem opposite to ours sometimes occurs in the modelling of communication problems, where two users wish to co-operate to avoid each other; either to transmit a message in a time slot in which the other is not also trying to transmit, or to choose a band to transmit on that is not already in use. In this context the Aloha and other protocols have been studied (see Abramson (1973), Tobagi (1980), Ross (1988), and Anatharam and Varaiya (1986)). In problems in which computers are communicating it is sometimes helpful to break the symmetry of the situation, by for example using an algorithm which incorporates a dependence on each computer's unique identifier. This type of procedure might also be helpful in the problem we are looking at, but we rule out this possibility by assuming that the two players are identical.

Some related problems have been considered by Schelling (1960) and Mosteller (1963). For example, Mosteller in his problem 'Quo Vadis' considered the case of two businessmen who wish to meet in New York City having neglected to name a meeting place. Mosteller's daughter said it was obvious that they should both go to the most famous place in the city. However, we are supposing that no such location exists.

A simple strategy consists of a decision on how to search amongst the different locations: this will be a choice of some sequence $\{s_1, s_2, s_3, \dots\}$, where each s_i is chosen from the set $\{1, 2, 3, \dots, n\}$. Since, before beginning this sequence, the player makes a random allocation of the labels $1, 2, 3, \dots, n$ to the n locations, if both players choose the same simple strategy they will meet if the two labellings coincide on one or more positions. However, there will be a positive probability that their labellings differ everywhere, in which case the two players will never meet. Since we wish to avoid this possibility it will be best to choose some mixed strategy, σ , giving a distribution over different simple strategies. An alternative way to describe such a mixed strategy is to specify for each of the finite sequences $\{s_1, s_2, \dots, s_k\}$ which might occur on the first k steps, the probabilities of choosing each of the n different possibilities on the next step.

Let H be the function giving the first meeting time of two fixed sequences of locations. Thus

$$H(\{s_1, s_2, \dots\}, \{t_1, t_2, \dots\}) = \min\{i : s_i = t_i\}.$$

When S is the sequence $\{s_1, s_2, s_3, \dots\}$ and ρ is a permutation we let ρS denote the sequence $\{\rho s_1, \rho s_2, \rho s_3, \dots\}$. Thus we may calculate the expected time, T , until two players meet each other, if both are using strategy σ , as

$$T_\sigma = E_{\pi_1, \pi_2, S_1, S_2} \{H(\pi_1 S_1, \pi_2 S_2)\}.$$

The expectation is over the random sequences S_1 and S_2 , chosen independently according to the strategy σ , and over π_1 and π_2 , which are independently and uniformly distributed over the $n!$ permutations of $\{1, 2, 3, \dots, n\}$ and correspond to the two labellings which the players adopt. The problem is to find a strategy σ which minimizes T_σ .

Since each player randomizes over all possible labellings, choosing a sequence S (say $S_\alpha = \{1, 2, 1, 3, 1, 2, 2, 1, \dots\}$ for example) will for any permutation ρ have precisely the same effect as choosing a sequence ρS (which might be $\{2, 1, 2, 3, 2, 1, 1, 2, \dots\}$ for the S_α above). Thus any mixture between such sequences will be equivalent. The consequence is that we may as well allow only strategies which have a 1 in the first position, a 2 as the next number appearing, a 3 as the next number which is neither 1 nor 2, and so on. More formally we say that a sequence $S = \{s_1, s_2, s_3, \dots\}$ (finite or infinite) is a *pattern* if for each i and j with $1 \leq i < j \leq n$,

$$\min\{k : s_k = i\} < \min\{k : s_k = j\}.$$

Without loss of generality we restrict attention to strategies which are described by a probability distribution over patterns.

It is sometimes helpful to keep a clear distinction between the actual locations and their labels. We continue to number the locations $1, 2, \dots, n$, but in the rest of this paper we use letters a, b, c, \dots etc. for their labels. Thus a pattern will have an a in the first position, and the first occurrence of b will be before the first occurrence of c , etc.

An obvious strategy to consider is the *random strategy* which at each state chooses each of the n different locations with probability $1/n$. This gives a value of T which is easy to calculate. Since $1/n$ is the probability that two players meet at the first step and if the two players do not meet on the first step then the whole process is repeated, we have

$$T = 1/n + (1 + T)(n - 1)/n.$$

Thus $T = n$. The question now arises as to whether we can achieve a better result than this by using some other strategy. We shall show that we can do better provided $n \geq 3$.

At this stage we need to make a comment concerning the problem formulation we have adopted, and its relationship to an actual problem of rendezvous between intelligent people. It is certainly true that since both players will adopt the same strategy the best that they can hope to achieve is an expected meeting time given by the solution to our problem. It is also true that if there is a single mixed strategy that is optimal for the problem then both players should clearly adopt it. If there are a number of different optimal strategies, however, it is far from clear how an intelligent person should decide between them. As we shall see in Section 3, this is exactly the situation that occurs when

$n = 3$. Furthermore, playing a mixture between different optimal strategies may not do very well. Rather than attempt a resolution of this difficulty, we suppose that both players are able to choose the same mixed strategy — perhaps they both carry with them a copy of this paper and are prepared to follow our instructions!

2. A search strategy for n locations

Suppose there are n locations. We shall show that the players can do substantially better than random search. That is, they can manage to meet in an expected time that is less than n . First, however, we state a lower bound that applies even after dropping the assumption that players must use the same strategy.

Lemma. *Suppose the players may agree to use different strategies. The only restriction is that they may not agree a common labelling of the locations. The time required to meet is at least $(n + 1)/2$ in expected value and it is stochastically minimized by a strategy which has one player remaining in one location and the other player searching locations in a random order without repetition.*

Proof. Let A_j be the event that the two players meet on the j th step (not necessarily for the first time). Because labels are applied randomly, $P(A_j) = 1/n$. Now the probability that they meet within k steps is

$$P\left(\bigcup_{j=1}^k A_j\right) \leq \min\left\{1, \sum_{j=1}^k P(A_j)\right\} = \min\{1, k/n\}.$$

The bound on the right-hand side is achieved by the strategy stated in the lemma. Clearly the expected meeting time under this strategy is given by $(1 + 2 + \cdots + n)/n = (n + 1)/2$.

Now we return to the assumption that the two players must use the same mixed strategy. Consider the following strategy. On the first step each player chooses a location at random; if they do not meet then their remaining strategy is described in terms of successive blocks of $n - 1$ steps. For the block of steps numbered $j(n - 1) + 2, \dots, j(n - 1) + n$, ($j = 0, 1, \dots$), each player does one of two things. Either he visits in some random order all of those $n - 1$ locations which differ from the location he last visited, or he remains in the location last visited for the whole block of $n - 1$ steps. He chooses between these options with probabilities $1 - \theta$ and θ respectively, independently of all previous random choices. The strategy renews every $n - 1$ steps.

We remark that we might have defined a similar strategy with blocks of n steps, where for each successive block of steps the strategy is to either remain in one randomly chosen location, or visit each of the n locations in some random order, with probabilities $1 - \theta$ and θ respectively. However, for any strategy that begins with a random choice of location the expected meeting time is $T = 1 + (n - 1)t/n$, where $(n - 1)/n$ is the probability of not meeting on the first step and t is the expected remaining time to meet following this strategy, conditional on not having met on the first step. Notice that if $t < n$ then $T > (t + (n - 1)t)/n$ and so $t < T$. But this would imply that if at any time

both players are about to disregard all previous history and restart from scratch by first searching a location at random (for an expected remaining meeting time of T) it will be better to pretend that this first random step has already occurred at the last step and obtain an expected remaining meeting time of t . We shall see below that indeed $t < n$ so it cannot be a good idea to disregard all previous history entirely. This suggests the rather more complicated strategy described above.

To analyse our strategy and determine the optimal choice of θ we need to know the probability with which two permutations of $\{1, 2, \dots, n\}$, chosen uniformly over the set of all $n!$ permutations, coincide in at least one place. Let $p(n, k)$ denote the probability that two such permutations meet in exactly k places. Now $p(n, k) = p(n - k, 0)/k!$ and $\sum_{k \geq 0} p(n, k) = 1$, so we have the formula

$$p(n, 0) = 1 - \sum_{k=1}^n p(n - k, 0)/k!$$

Given that two randomly chosen permutations meet in k locations, the probability that the first such meeting occurs no sooner than in the j th location is

$$q_j = \binom{n+1-j}{k} / \binom{n}{k}, \quad j = 1, \dots, n+1-k.$$

Hence the expected number of steps until the first meeting occurs is $q_1 + \dots + q_{n+1-k}$, which turns out to be just $(n+1)/(k+1)$. If we write τ for the time at which the first meeting occurs in two randomly chosen tours of n locations, then

$$E\{\min(\tau, n)\} = p(n, 0)n + \sum_{k=1}^n p(n, k) \left(\frac{n+1}{k+1}\right).$$

It is not hard to see that this expression can be rewritten as $(n+1)\{1 - p(n+1, 0)\} - p(n, 0)$. Now notice that $E\{\min(\tau, n)\} = 1 + (1 - 1/n)E\{\min(\tau, n) \mid \tau > 1\}$. Thus, given that meeting does not occur at step 1, the remaining expected time to meet, say t , is

$$t = \theta^2(n - 1 + t) + 2\theta(1 - \theta)(n/2)$$

$$(1) \quad + (1 - \theta)^2 \frac{n}{n-1} \{p(n, 0)t + (n+1)(1 - p(n+1, 0)) - p(n, 0) - 1\}$$

$$= \frac{\theta^2(n-1) + \theta(1-\theta)n + (1-\theta)^2 \frac{n}{n-1} \{(n+1)(1 - p(n+1, 0)) - p(n, 0) - 1\}}{1 - \theta^2 - (1-\theta)^2 \frac{n}{n-1} p(n, 0)}.$$

The expression (1) can be minimized over θ . The expected meeting time, accounting for the first random step is then $T = 1 + (n-1)t/n$. Some values are given in Table 1.

Note that asymptotically the optimal value of θ tends to a limit 0.24748942. This is found from (1) by using the well-known result that as $n \rightarrow \infty$, $p(n, 0) \rightarrow 1/e$. The expected meeting time is asymptotically $(0.82888497)n$. (We have stated so many significant

TABLE 1

n	Optimal θ	T	T/n
2	0.5000	2.0000	1.0000
3	0.3333	2.6667	0.8889
4	0.3220	3.5685	0.8921
5	0.3012	4.3793	0.8759
6	0.2914	5.2133	0.8689
7	0.2842	6.0421	0.8632
8	0.2791	6.8716	0.8590
9	0.2753	7.7008	0.8557
10	0.2722	8.5300	0.8530
50	0.2521	41.6871	0.8338
$n \rightarrow \infty$	0.2475		0.8289

figures here in order that anyone trying to find an asymptotically better strategy will know if they have succeeded).

The strategy described in this section roughly corresponds to the behaviour of people faced with this type of problem in practice. They spend some of their time stationary, hoping the other person will come to find them, and they spend some of their time moving about looking for the other person. Obviously it is bad for both people to be stationary, and intuitively one feels that it is not a good idea for both to be moving. If only they had agreed beforehand who would stay put and who would search, then they would meet in $(n + 1)/2$ steps on average (see the lemma at the start of this section). Lacking this agreement, we recommend they each review the situation every $n - 1$ steps and then stay where they are or set out on a search with probabilities of about 0.25 and 0.75.

3. Optimal strategies for two and three locations

In the previous section we described a form of strategy in which over blocks of $n - 1$ steps each player either remains in one location for $n - 1$ steps or visits $n - 1$ different locations. For $n = 2$ this is simply the random strategy of visiting locations randomly at each step. For $n = 3$ the situation is already much more complicated. However before beginning our discussion of the case of three locations we demonstrate that the random strategy is optimal if there are only two locations.

Theorem 1. For $n = 2$ the minimal expected time to rendezvous is 2. This is achieved by the random strategy, which stochastically minimizes the time to rendezvous.

Proof. If the two players have not met over the first k steps then they must have chosen sequences which were exactly opposite, and thus which were generated with different labels from the same pattern, say pattern P . Let p be the probability with which under an optimal strategy the pattern P is followed by an a on the $(k + 1)$ th step. Since both players use the same strategy the probability that they meet at the $(k + 1)$ th step is $2p(1 - p)$, which has a maximum value of $1/2$. Thus the probability of meeting at any

step is at most $1/2$, independently of the pattern that has been followed so far. Since the random strategy achieves this on every step it stochastically minimizes the expected number of steps until successful rendezvous.

For $n = 3$ the strategy described in Section 2 is more interesting. It can be characterized as a strategy in which over the first three steps each player uses patterns aaa or abc with probabilities $1/3$ and $2/3$ respectively. If a meeting has not occurred following step j , ($j = 3, 5, 7, \dots$) then for the next two steps a player either remains where he is, or visits in random order the two locations that differ from the location he was in at step j , with probabilities $1/3$ and $2/3$ respectively. The strategy has an expected meeting time of $8/3$. We shall show that this strategy is optimal.

Observe that an equivalent strategy is one in which, if a meeting has not occurred following step j , ($j = 1, 3, 5, 7, \dots$), then for the next two steps a player either returns to the location he visited at step 1, or he visits in random order the two locations that differ from the location he visited at step 1, with probabilities $1/3$ and $2/3$ respectively. Call this strategy π . In claiming that it is an equivalent strategy we have observed that it does not matter which location is given special status as the location in which a player will stay for steps $j + 1$ and $j + 2$ with probability $1/3$, provided both players choose it in the same way. Strategy π chooses the pattern for steps $j + 1$ and $j + 2$ as equally likely to be aa , bc or cb , independently of all previous history but the first step.

In order to prove that π is optimal we shall show that even after giving the players some additional information the expected meeting time is still at least $8/3$. The idea of giving extra information and then bounding the performance of an optimal strategy using this information has been used to good effect in some other unusual problems (see Anatharam and Varaiya (1986)). The extra information we supply is as follows. Imagine the three locations as equally spaced around a circle. Suppose the first step is taken and the players do not meet. Suppose now that a direction is randomly imposed on the circle, clockwise or counterclockwise with equal likelihood, and that the choice of this direction is communicated to both players. Each player now has an unambiguous method of referencing the locations: he can let a be the location he visited at step 1, and let b and c be respectively the locations that are next on and second on from location a in the direction that has been imposed on the circle. Clearly, since both players can disregard the additional information if they choose, it can only decrease the expected remaining time to meet. In so far as it affects the remaining expected time until they meet, the circumstances with this additional information are identical to those that would arise if both players were to use the pattern abc over the first three steps and they were not to meet. This follows because if each player chooses three different locations and they do not meet, then the two sequences of locations must be cyclic permutations of each other. More specifically if one player visited locations in order 1, 2, 3 then the other must have visited the locations in order 2, 3, 1 or 3, 1, 2.

Theorem 2. For $n = 3$ the minimal expected time to rendezvous is $8/3$. This is achieved by strategy π .

Proof. From the remarks above, the expected time to meeting is at least $1 + (2/3)t$, where t is the expected remaining time to meet given that each player has on the first

three steps visited all three locations and they have not met. Thus it suffices to show $t \geq 5/2$. So suppose both players have followed pattern abc for three steps and they have not met. Let p_{ij} be the probability with which on the next two steps a player first visits the location that was label i in his initial triple and then visits the location that was label j in his initial triple, $i, j \in \{a, b, c\}$. Let $p_i = \sum_j p_{ij}$. The expected time to meet is at least T , where

$$(2) \quad T \geq \sum_i p_i^2 \left(1 + T + \sum_j p_{ij}^2 / p_i^2 \right) + (1/2) \sum_{\substack{ij,kl \\ i \neq k}} p_{ij} p_{kl} A_{ij,kl}$$

and $A_{ij,kl}$ is given below, with its rows and columns arranged to correspond to the i, j pairs ($aa, ab, ac, bb, bc, ba, cc, ca, cb$) in that order.

$$A = \begin{bmatrix} 9 & 7 & 7 & 4 & 3 & 4 & 4 & 4 & 3 \\ 7 & 9 & 7 & 4 & 4 & 3 & 3 & 4 & 4 \\ 7 & 7 & 9 & 3 & 4 & 4 & 4 & 3 & 4 \\ 4 & 4 & 3 & 9 & 7 & 7 & 4 & 3 & 4 \\ 3 & 4 & 4 & 7 & 9 & 7 & 4 & 4 & 3 \\ 4 & 3 & 4 & 7 & 7 & 9 & 3 & 4 & 4 \\ 4 & 3 & 4 & 4 & 4 & 3 & 9 & 7 & 7 \\ 4 & 4 & 3 & 3 & 4 & 4 & 7 & 9 & 7 \\ 3 & 4 & 4 & 4 & 3 & 4 & 7 & 7 & 9 \end{bmatrix}$$

The explanation of (2) is as follows. If on the fourth step the players both choose their i th symbol, which occurs with probability p_i^2 , then the total expected time must be at least $1 + T$. If, having done this, they then both choose their j th symbol for the fifth step, which occurs with probability p_{ij}^2 , then the total expected time is at least $2 + T$. This explains the first term on the right-hand side in (2). If, however, their patterns differ as a result of the fourth step, then with probability $1/2$ they meet on this step. For example, if the first player who has visited locations 1, 2, 3 on his first three steps chooses a on his next step, therefore visiting location 1, and the second player, who is equally likely to have visited 2, 3, 1 or 3, 1, 2 chooses b on his next step, then this second player is equally likely to choose location 3 or 1. On the other hand if the second player chooses c , he is equally likely to choose location 1 or 2. If the patterns for the two players differ on the fourth step and they do not meet, then it is straightforward to check that they meet on the following (fifth) step only if their choices of i, j are distinct pairs from the same one of the three sets: $\{aa, bc, cb\}$, $\{ac, bb, ca\}$ and $\{ab, ba, cc\}$. Otherwise they will require at least three steps to meet. So, for example, if one player does aa and the other ba , then the expected meeting time is at least $(1/2)1 + (1/2)3 = 2$. On the other hand, if one player does aa and the other bc , then the expected meeting time is exactly $(1/2)1 + (1/2)2 = 3/2$. Note that the probabilities of these events are $2p_{aa} p_{ba}$ and $2p_{aa} p_{bc}$ respectively.

These bounds give rise to the 4s and 3s in the matrix A and explain the final term on the right-hand side of (2).

We wish to prove $T \geq 5/2$. From (2) we obtain a lower bound for T , say LB , where

$$LB = \inf_{p_i} \left\{ \sum_i p_i^2 + \sum_{i,j} p_{ij}^2 + (1/2) \sum_{\substack{i,j,k,l \\ i \neq k}} p_{ij} p_{kl} A_{ij,kl} \right\} / \left\{ 1 - \sum_i p_i^2 \right\}.$$

It follows that if $5/2$ exceeds LB , then we would have

$$(3) \quad \begin{aligned} 5/2 &> \sum_i p_i^2 \left(1 + 5/2 + \sum_j p_{ij}^2 / p_i^2 \right) + (1/2) \sum_{\substack{i,j,k,l \\ i \neq k}} p_{ij} p_{kl} A_{ij,kl} \\ &= (1/2) \sum_{i,j,k,l} p_{ij} p_{kl} A_{ij,kl}. \end{aligned}$$

The 7s and 9s in A come from the constants under the first two of the three summations in the right-hand side of (3). We wish to show that (3) cannot hold, or in other words that the quadratic form, $\sum_{i,j,k,l} p_{ij} p_{kl} A_{ij,kl}$, cannot take a value strictly less than 5 for any probabilities p_{ij} .

Consider now the quadratic form, $p^T A p$. For strategy π , which we have conjectured to be optimal, $p^T = (1, 0, 0, 0, 1, 0, 0, 0, 1)/3$ and $(p^T A)_i = 5$. Thus if $q = p + x$ is any other probability vector, so that $x^T 1 = 0$, we have $(p + x)^T A (p + x) = p^T A p + 2p^T A x + x^T A x = 5 + 0 + x^T A x$. Now A is positive semi-definite (with eigenvalues 0, 0, 3, 3, 3, 3, 12, 12 and 45) so we have $q^T A q \geq 5$ and no strategy is better than π .

Interestingly, π is a one-step-look-ahead strategy. That is, if it has been followed for k steps then following it for one more step maximizes the probability of rendezvous at the next step. However, we do not expect the optimal strategy to be one-step-look-ahead in general. For $n = 4$ one can show that a one-step-look-ahead strategy follows patterns $aaaa$ and $abcd$ over the first four steps with probabilities $1/4$ and $3/4$ respectively. These are not the probabilities which arise using the strategy for $n = 4$ described in Section 2.

4. Extensions and open problems

An attractive feature of this problem is that there are a great many related problem formulations which are of interest. In this section we review a number of these.

4.1. *Locations with structure.* One idea is to introduce some structure in the set of locations. This is most conveniently done by supposing that there are connections between some pairs of locations, for example we might suppose that all the locations are connected in a loop (perhaps with directed arcs). We wish to retain the idea that the two players are unable to share the same labelling, so we would not want to give a structure which enables the two players to identify one of the locations as different. Thus for example a star structure, where one location is connected to all the others, would suggest that both players should go to that location immediately. It is possible,

however, that an arrangement with the locations arranged in a loop with directed arcs around the loop, does not reduce the expected time till the players meet. This is certainly true for three locations and indeed we used precisely this fact in proving the optimality of the strategy π . It is an open question whether the same thing is true with four or more locations.

If the locations have directed arcs between them it is natural to think of a problem in which the players may only move along arcs, so that at each step they either stay in the same place or move to an adjacent location. One very simple problem in this class has three locations joined in a circle with directed arcs, so that a player at each stage has the choice of whether to stay in the same location or move around to the next one. The analysis of this problem turns out to be very similar to that in which we can move arbitrarily amongst three locations, and we can establish the optimal strategy using the same method as in the proof of Theorem 2. The optimal strategy is to choose the first location at random and to then decide subsequent moves in pairs (as for strategy π). After each odd numbered step a player should with equal probabilities either stay in his present location for the next two steps, or move on one location at each of the next two steps. The expected time to meet is $10/3$.

If arcs are undirected it will be necessary to specify whether two players moving in opposite directions on the same arc will meet each other. Suppose that two players cannot pass each other on the same arc without meeting and consider the case in which there are a very large number of locations arranged in a loop. It is easy to see that this problem is essentially equivalent to that of attempting to rendezvous on a circle, when neither player knows the starting position of the other, nor the orientation of the circle that the other is using. Steve Alpern has, in conversation, conjectured that for this problem the optimal strategy is for both players to make successive half tours around the circle in randomly chosen directions (the so called 'cohato' strategy, see Foreman (1977)).

4.2. Variations on the search model. The problem we have considered makes the assumption that each player searches just one location and that if the players are both searching the same location they will certainly find each other. These assumptions can be dropped. Most straightforwardly we may suppose that the probability of the two players finding each other if they both search the same location is given by some fixed probability α , which is the same for each location.

Here $E(T) = \sum_{k \geq 1} \alpha(1 - \alpha)^{k-1} E(T_k)$, where T_k is the number of the step on which they are in the same location for the k th time. When $n = 2$ the random strategy is still optimal. This follows from the observation that this problem is equivalent to our original problem, but with the added factor that certain periods have been (secretly) marked as ones during which meeting cannot occur even when players occupy the same location. Each period has a probability $1 - \alpha$ of being marked, so that the marked periods are distributed as a renewal process, separated by geometrically distributed numbers of unmarked periods. If the players were told which periods had been marked, then it would be optimal (by Theorem 1) for them to follow the random strategy during unmarked periods. They would be indifferent as to what they did during marked periods,

so they might as well follow the random strategy. This means they can optimally follow the random strategy for every period, regardless of which periods are marked.

Another alternative is to suppose that each player is able to search m of the n locations at each stage. The simplest problem of interest in this case has $m = 2$ and $n = 4$. For this problem if the two players miss each other up to time k , then they must have chosen mirror image pairs of locations at each stage. It is easy to show, using the same technique as in the proof of Theorem 1, that the optimal strategy is to choose at each stage each of the six possible pairs of different locations (i.e. ab , ac , ad , bc , bd , cd) with probability $1/6$. It is an open problem to find the optimal strategy when $m = 2$ and $n = 5$.

Yet another variation on the problem, posed when there is a network structure on the locations, is one in which the two players meet whenever they search either the same location or adjacent locations. For certain network structures this is equivalent to the problem in which m out of n locations can be searched. Thus, for example, the $m = 2$ and $n = 4$ problem is equivalent to the rendezvous problem on six locations at the vertices of an octahedron (where the two players meet if they are at either the same or adjacent vertices).

4.3. Questions of timing. An important element in all the problems considered so far is the synchronization of moves by the two players. This simplifies the analysis, but might be regarded as very unrealistic in terms of the original problem in which two people are attempting to meet in a building. It is possible to consider a weakening of this synchronization in two different ways. First we might suppose that the two players start searching at different times, with neither being aware of whether the other is already searching when he arrives at the first location. In this situation we can continue to assume that the two players move between locations at the same instant. Then we would be concerned to minimize the expected time till they meet from the arrival of the second player. The optimal strategy in this case is certain to be very complex. Even if we require each player to use a strategy which does not depend on the time that he starts searching, the optimal strategy might still depend on what the players know about the distribution of the difference in their starting times. It could be that a strategy which is optimal for synchronized starting times now does worse than the random strategy. However, when $n = 3$ our strategy π still does better than the random strategy. This follows from the two-step nature of strategy π (which has players choosing at random between patterns aa , bc or cb for each pair of successive steps). If the second player to start searching does so at a time that results in the players being synchronized in their applications of π then the expected time until they meet is $8/3$ as before. If, however, the second player starts at a time that makes their applications of π unsynchronized, then at each step at least one of the players will select his next location randomly; so there will be a $1/3$ probability of meeting in each period and the expected time until rendezvous will be 3.

Another possibility is to allow the players to move between locations at any time. In this case we wish to make the probability of meeting, when both players are searching the same location, depend on the amount of time that they have been there. This problem is similar to that we mentioned above, where the probability of finding the other player on one step is reduced to α when the two players are in the same location. The

continuous-time version can be obtained as the limit of discrete-time versions as α approaches zero.

4.4. *Leaving messages.* Shmuel Gal has suggested a problem in which each player may leave messages in the locations he visits. What should the messages say if the two players wish to meet and how quickly can meeting occur?

Clearly the expected time until at least one player first finds a message left by the other player is minimized by each player visiting locations at random without returning to any locations previously visited. The probability that no location has been visited by both players within the first k steps is then $(1 - k/n)(1 - k/(n - 1)) \cdots (1 - k/(n - k + 1))$. So the expected time, S , until one first finds a message left by the other is

$$S = \sum_{k=1}^{\lfloor n/2 \rfloor} \left\{ \prod_{j=1}^k (1 - k/(n - j + 1)) \right\}.$$

It is interesting to see that S/\sqrt{n} tends to $(1/2)\sqrt{\pi}$. An upper bound for S is given by

$$S \leq \sum_{k=1}^{\infty} (1 - k/n)^k \leq \sum_{k=1}^{\infty} \exp(-k^2/n) \leq \int_{x=0}^{\infty} \exp(-x^2/n) dx = (1/2)\sqrt{n\pi}.$$

To obtain a lower bound we use an argument supplied by Leo Flatto and Andrew Odulyzko of AT&T Bell Laboratories. Let θ be chosen $1/2 < \theta < 2/3$. Then for $1 \leq k \leq n^\theta$ and $j < k$, we have

$$\begin{aligned} &\log(1 - k/(n - j + 1)) \\ &\geq \log(1 - k/(n - k)) = -k/(n - k) + O(k^2/(n - k)^2) = -k/n + O(k^2/n^2). \end{aligned}$$

So

$$\prod_{j=1}^k (1 - k/(n - j + 1)) = \exp(-k^2/n + O(k^3/n^2)) = \exp(-k^2/n + O(n^{3\theta-2}))$$

whence

$$\begin{aligned} S/\sqrt{n} &\geq \left\{ \sum_{k=1}^{n^\theta} (1/\sqrt{n}) \exp(-k^2/n + O(n^{3\theta-2})) \right\} \simeq \int_{x=0}^{\infty} (1/\sqrt{n}) \exp(-x^2/n) dx \\ &= (1/2)\sqrt{\pi}. \end{aligned}$$

Thus we see that, no matter what sort of messages the players use, the expected time required to rendezvous is at least $O(\sqrt{n})$. However, even when players leave only the simplest messages, of the form 'I was here', the following algorithm achieves meeting in a time which is $O(\sqrt{n})$ on average.

Step 1. For exactly \sqrt{n} steps visit \sqrt{n} different locations and leave messages. Call these \sqrt{n} locations your *original set*. Disregard for the moment any messages you find left by the other player.

Step 2. For the next \sqrt{n} steps revisit each member of your original set once and note at which locations the other player has also left messages.

Step 3. If Step 2 discovered a non-empty intersection between your original set and

that of the other player you should now continue to revisit locations in this intersection, at random, until you meet the other player. Otherwise continue to Step 4.

Step 4. Visit at random locations not yet visited. Do not leave messages. Continue until you first visit a location where a message has been left by the other player. Leave a message in this location.

Step 5. Revisit locations in your original set at random until you discover a message left by the other player. From now on distribute your visits uniformly between the two locations which you know you have both visited, until you meet.

If the algorithm stops at Step 3 then the expected meeting time is at most $3\sqrt{n}$. If it continues to Step 5, the expected meeting time is at most $6\sqrt{n} + 2$.

5. Conjectures

One tantalizing aspect of the rendezvous problem is that the proofs of even simple results seem to be elusive and require new techniques. A proof of the optimality of strategy π for $n = 3$ was more difficult than one might have expected. We conclude the paper by making a number of conjectures. Even conjecture (1), which at first sight seems almost immediate, appears not to be straightforward to establish.

(1) The minimal expected number of steps required to meet is an increasing function of n .

(2) For $n = 3$, strategy π maximizes for all k the probability of meeting within k steps. (We know this is true for $k \leq 4$.)

(3) The policy described in Section 2 is not optimal for $n = 4$.

(4) As $n \rightarrow \infty$ the minimal value of $T/n \rightarrow 0.8289$. That is, the policy described in Section 2 is an asymptotically optimal strategy.

(5) In the case that players who search the same location only meet with a probability α , we have shown that the optimal strategy for $n = 2$ is unchanged, regardless of the value of α . We conjecture that even for larger n the optimal strategy is the same for all α .

References

- ABRAMSON, N. (1973) Packet switching with satellites. *Proc. NCC*, 695–702.
- ALPERN, S. AND ASIC, M. (1986) Ambush strategies in search games on graphs. *SIAM J. Control Optim.* **24**, 66–75.
- ANANTHARAM, V. AND VARAIYA, P. (1986) An optimal strategy for a conflict resolution problem. *Syst. Control Lett.* **6**, 329–332.
- FOREMAN, J. G. (1977) Differential search games with mobile hider. *SIAM J. Control Optim.* **15**, 841–856.
- GAL, S. (1980) *Search Games*. Academic Press, San Francisco.
- MOSTELLER, F. (1963) *Fifty Challenging Problems in Probability with Solutions*. Addison-Wesley, New York.
- ROSS, S. M. (1988) A simple proof of instability of a random-access communication channel. *Prob. Eng. Inf. Sci.* **2**, 383–385.
- SHELLING, T. C. (1960) *The Strategy of Conflict*. Oxford University Press.
- TOBAGI, F. A. (1980) Multi-access protocols in packet-communications systems. *IEEE Trans. Commun.* **28**, 468–488.
- WEBER, R. R. (1986) Optimal search for a randomly moving object. *J. Appl. Prob.* **23**, 708–717.