This column is a place for those bits of contagious mathematics that travel from person to person in the community, because they are so elegant, suprising, or appealing that one has an urge to pass them on.

Contributions are most welcome.

Please send all submissions to the Mathematical Entertainments Editor, **Ravi Vakil,** Stanford University, Department of Mathematics, Bldg. 380, Stanford, CA 94305-2125, USA e-mail: vakil@math.stanford.edu

The Locker Puzzle

Eugene Curtin and Max Warshauer

Suppose I take the wallets from you and ninety-nine of your closest friends. We play the following game with them: I randomly place the wallets inside one hundred lockers in a locker room, one wallet in each locker, and then I let you and your friends inside, one at a time. Each of you is allowed to open and look inside of up to fifty of the lockers. You may inspect the wallets you find there, even checking the driver's license to see whose it is, in an attempt to find your wallet. Whether you succeed or not, you leave all hundred wallets exactly where you found them, and leave all hundred lockers closed, just as they were when you entered the room. You exit through a different door, and never communicate in any way with the other people waiting to enter the room. Your team of 100 players wins only if *every team member* finds his or her own wallet. If you discuss your strategy beforehand, can you win with a probability that isn't vanishingly small?

We develop a more mathematical formulation to facilitate a precise discussion of the problem. This consists of numbering our players, and replacing wallets by player numbers! Our game is played between a single Player A against a Team B with 100 members, $B_1, B_2, \ldots, B_{100}$. Player A places the numbers $1, 2, \ldots, 100$ randomly in lockers $1, 2, \ldots, 100$ with one number per locker. The members of Team B are admitted to the locker room one at a time. Each team member is allowed to open and examine the contents of exactly 50 lockers. Team B wins if every team member discovers the locker containing his own number. Team B is allowed an initial strategy meeting. No communication is allowed after the initial meeting, and each team member must leave the locker room exactly as he found it. It is important to realize that the solution does not involve some trick to pass information from one player to another. We could equally well make 100 copies of the room and make an identical distribution of numbers into lockers for each room, then ask the members of Team B to perform their searches simultaneously, with one person per room.

Each individual will succeed in finding his own number with probability $^1/_2$. If they act independently, they must get lucky 100 times in a row, and the team will win with probability only $(^1/_2)^{100}$. Team B needs some help! Amazingly there is a strategy which gives significant probability of success for Team B. Even if we give the problem with $2n$ players on Team B each of whom can examine n out of $2n$ lockers, Team B can apply the strategy to succeed with probability over 30% regardless of how large a value we take for n. Your problem is to find this strategy.

Searching For Ideas

Let's play with some ideas using a more manageable number of players. To be as concrete as possible, let's switch to the case of 10 players on Team B, each of whom can examine 5 out of 10 lockers. Here random guessing by each player is already somewhat hopeless and succeeds with probability $(^1/_2)^{10} = \frac{1}{1024}$. A first try to improve the probability of success is to search for a clever way to assign a set of lockers for each person to examine. Certainly we can improve over random guessing in this manner. For example if team members 1–5 examined lockers 1–5, and team members 6–10 examined lockers 6–10, they would succeed provided numbers 1–5 are placed in lockers 1–5. Number 1 is placed somewhere in the first 5 lockers with probability 5/10, then given that number 1 is so placed, number 2 is also in the first 5 with probability 4/9 and so on. Following this plan, Team B will succeed with probability

$$\frac{5}{10}\,\frac{4}{9}\,\frac{3}{8}\,\frac{2}{7}\,\frac{1}{6} = \frac{1}{242}.$$

While this is an improvement over random guessing, it still leaves Team B with slim chances. Although the scheme fails, it is worth noticing that if B_1 finds his number in this scheme, then B_6 will find his number with probability 5/9 (as he will look in 5 lockers not including the one containing the number 1), but B_2 will find his with probability only 4/9. The success or failure of B_1 can influence the probabilities of success of the other members. This is the first clue!

An ideal strategy would be one where if B_1 succeeds then everyone else does too. Note that this would allow the whole team to succeed half the time even though each individual member fails half the time. This ideal is not attainable, but perhaps you can find a strategy where if B_1 succeeds, then everyone else is more likely to succeed. No method of preassigning lockers will accomplish this, as if B_1 finds his number in locker k anyone with locker k in their preassigned set has his chances reduced. This suggests that the locker choices will have to depend on information not available at the initial meeting. The only such information available is the numbers a player finds inside the lockers he opens. With this further hint try one more time to find a good strategy before we proceed to the solution!

Developing the Solution

Once we realize that the locker B_1 opens at any stage can depend on what he has found inside the lockers he has already opened, the number of possible strategies to consider is enormous, even in the 10-player case. The strategy must tell B_1 which locker to open first (10 choices), which locker to open next if he is not lucky on the first try (9 choices for each of the possible 9 numbers he may see), which to open third if he is not lucky on his second attempt either (8 choices for each of the 9×8 possible sequences of 2 numbers he has seen so far), and so on. So B_1 alone has $10 \times 9^9 \times 8^{9\times 8} \times 7^{9\times 8\times 7} \times 6^{9\times 8\times 7\times 6}$ possible strategies. To compute the number of strategy choices for the whole team, we raise this to the 10th power and get a number 28,537 digits long! How are we to choose one?

In this section we will show that one very simple strategy lets the team win with remarkably high probability. The strategy for any one player is entirely unremarkable; the magic arises from the fact that the chances of the different players winning are highly correlated. Moreover, in the next section, we will show that the strategy is in fact optimal.

Fortunately the good strategy is simple to implement and the choice of the next locker does not depend on the entire sequence of numbers seen but only on the most recent number. The good strategy has player B_i start by opening locker i. Then if he finds number k at any stage and $k \neq i$, he opens locker k next. Notice that player B_i, never opens a locker (other than locker i) without first finding its number, so each time he opens a new locker he must find either his own number or the number of another unopened locker.

Again let's look at a particular case with 10 players and suppose, for example, that the numbers are distributed in the order 6,8,9,7,2,4,1,5,10,3. Player B_1 first examines locker 1 and finds the number 6. So he looks in locker 6 finding the number 4, then locker 4 finding the number 7, then finally in locker 7 finding his number. When he finds his number, B_1 will now know that B_6, B_4, and B_7 will look in exactly the same lockers in the same cyclic order, each finding his number on the 4th try! He also knows that none of the other players will waste any tries on these lockers.

We may represent any permutation of numbers into lockers by listing the cycles. The permutation 6,8,9,7,2,4,1,5,10,3 gives the cycles (6, 4, 7, 1) (8, 5, 2) (9, 10, 3), and Team B succeeds because there is no long cycle. To find the probability that Team B wins, we count the number of permutations of 10 numbers with a cycle of length 6 or longer. First let's count how many have a 6-cycle. Choose which 6 elements go into the 6-cycle, arrange them in cyclic order, and then pick an arbitrary permutation of the remaining 4 elements. The number of ways to do this is

$$\binom{10}{6} 5!4! = \frac{10!5!4!}{6!4!} = \frac{10!}{6}$$

So 1/6 of the 10! permutations have a 6-cycle, and a random permutation has a 6-cycle with probability 1/6. The same argument can be used to find the probability of a permutation of 1–10 having a cycle of any length longer than 6. (We warn that the argument does not work for counting the number of permutations of 1–10 with a 5-cycle (or shorter) as the permutation could have two 5-cycles.) A permutation of 10 numbers has a 7-cycle with probability 1/7 and so on, and the probability of a cycle of length 6 or larger is $1/6 + 1/7 + 1/8 + 1/9 + 1/10 = 1627/2520 \approx 0.645635$. This gives the probability that Team B will fail, so of course Team B wins with probability $1 - 1627/2520 = 893/2520 \approx 0.354365$. Over 35% of the time, all 10 members of Team B find their own wallets!

Will this idea be good enough for the initial version with 100 players? We can do the analogous computation and see that this pointer-following strategy works with probability $1 - (1/51 + 1/52 + \cdots + 1/100) \approx .311828$.

Notice that while our strategy has still performed remarkably well for 100 players, the probability of success was still less than in the 10-player version. As we increase the number of players, does the success rate decrease to zero, or does it always stay above a certain positive number? With $2n$ players and $2n$ lockers, Team B will win provided that the permutation of numbers in lockers has no cycle of length $n + 1$ or longer. The probability of such a long cycle is $\Sigma_{k=1}^{n} \frac{1}{n+k}$. By viewing this expression as an upper Riemann sum for $\int_n^{2n} \frac{1}{x+1}\,dx$ and a lower Riemann sum for $\int_n^{2n} \frac{1}{x}\,dx$ we obtain

$$\ln\left(2 - \frac{1}{n+1}\right) = \int_n^{2n} \frac{1}{x+1}\,dx \le \sum_{k=1}^{n} \frac{1}{n+k} \le \int_n^{2n} \frac{1}{x}\,dx = \ln 2.$$

So $\Sigma_{k=1}^{n} \frac{1}{n+k} \to \ln 2$ as $n \to \infty$; moreover the sum increases monotonically with n. So the expression $1 - \Sigma_{k=1}^{n} \frac{1}{n+k}$ giving the probability of success is monotonically decreasing to $1 - \ln 2 \simeq 0.306853$. Team B wins with the pointer-following strategy with probability exceeding 30%, regardless of the number of players and lockers. Now that we have found a good strategy, we turn our attention to whether it provides the best possible solution.

Is Pointer-Following Optimal?

We establish the optimality of pointer-following by comparing the game considered above (Game 1) with a new game (Game 2) between the same adversaries, Player A and Team B. For simplicity we give the argument in terms of the 10-player versions. Recall that in Game 1 we are allowing each player to examine 5 lockers. We first modify this rule and say that each player must continue examining lockers until he has opened the locker containing his number, and then he is not allowed to open any further lockers. Team B wins if no player opens more than 5 lockers. This change makes no difference to who wins in Game 1, but it will clarify the comparison with Game 2.

In Game 2, Player A again distributes the 10 numbers at random in the 10 lockers. Then all of team B is invited into the locker room together. Team member B_1 is required by the rules to start opening lockers and continue until she reveals the number 1. Once she has opened the locker containing the number 1, she may not open any further lockers; then, the lowest-numbered member of Team B whose number has not yet been revealed is required to take over opening lockers until he finds his number and so on. Team B continues until all lockers are opened. Again Team B wins if no individual team member opens more than 5 lockers. Before proceeding, we invite you to consider the following questions: With what probability can Team B win Game 2? What strategy should the team members employ? Does their choice of strategy even matter?

Let's sit in the locker room and observe Team B in the process of playing Game 2. We record the progress, listing the numbers in the order in which they are revealed. Our list of numbers is sufficient to determine how many lockers were opened by each player. For example, if we record the list 2,6,1, 4,9,7,10,8,3,5, we know that player B_1 revealed the numbers 2, 6, and 1. Then player B_3 was required to take over, and he opened the lockers containing the numbers 4, 9, 7, 10, 8, and 3, in that order. Then player B_5 opened the remaining locker containing the number 5. In this example Team B lost, as player B_3 opened 6 lockers. Notice that we will record any given ordering of the numbers 1–10 with probability 1/10!. The first number revealed is 2 with probability 1/10, no matter which locker is opened, given that the first is 2 the second will be 6 with probability 1/9, and so on. What strategy is Team B following here? It makes absolutely no difference! Team B can choose lockers at random or follow the most sophisticated plan; we still get probability 1/10! for each of the 10! possible orders in which the numbers could be revealed. In Game 2 Team B's probability of success is completely independent of strategy.

To find the probability that Team B wins, we must count how many of the 10! possible orders of the numbers 1–10 represent wins. We employ a version of the classical records-to-cycles bijection [6, p17] to assign a permutation written in cycle notation to each ordering. The first cycle of our permutation consists of the numbers opened by B_1 in order; the second cycle, the numbers opened by the second locker opener; and so on. So, for example, 2,6, 1,4,9,7,10,8,3,5 $\rightarrow$ (2,6,1)(4,9,7,10,8,3)(5). Furthermore we see that each permutation arises in this manner from a unique ordering of the numbers 1–10. We first write the permutation in cycle notation, rotate each cycle so that the lowest number in the cycle is written last, and then order the cycles so that their last numbers are in ascending order. For example (9,7,8)(1,3,10,5)(2,4,6) = (3,10,5,1)(4,6,2)(8,9,7) $\rightarrow$ 3,10,5,1,4,6,2,8,9,7. We have established a one-to-one correspondence between lists for which Team B wins and the permutations of 1–10 with no cycles of length greater than 5. Thus the probability that Team B wins Game 2 is the probability that a random permutation of 1–10 has no cycle of length greater than 5, and we have already computed this as $893/2520 \approx 0.354365$. This is exactly the probability of success for Team B in Game 1 using pointer following!

Our analysis has a significant consequence for Game 1. Team B can take any Game 1 strategy and adapt it to Game 2 as follows: If player B_i is opening lockers in Game 2, he can use his Game 1 strategy for choosing lockers to open, simply observing the contents without wasting a turn if the indicated locker is already open. Thus if a strategy succeeds in Game 1 for a particular distribution of numbers into lockers it will also succeed in Game 2. If there were a better strategy for Game 1 we could apply it in Game 2 and get a better chance to win this game also. But this is impossible, as all strategies for Game 2 lead to the same probability of success.

We have one final small puzzle: Happy with their optimal strategy for Game 1, Team B began a sequence of matches with Player A, but they soon found themselves down 10 to 0. What do you suspect Player A is doing? (It seems that Player A subscribes to the *Intelligencer* and has devised a plan to defeat Team B.) What can Team B do to counter Player A's plan?

History of The Locker Puzzle

Our problem was initially considered by Peter Bro Miltersen, and it appeared in his paper [4] with Anna Gál, which won a best paper award at the ICALP conference in 2003. Miltersen says of the problem, "I think it started spreading when I presented it to several people at Complexity 2003, which was held in Aarhus, where I was a local organizer." In their version there is one numbered slip of paper for each player on the team. Player A then colors each slip either red or blue. Each member of Team B may examine up to half the lockers. He is then required to state or guess the color of the slip of paper with his number. Again every team member must state or guess his color correctly for the team to win. Initially Miltersen expected that Team B's probability of success would approach zero rapidly as the number of players increased. However, Sven Skyum, a colleague of Miltersen's at the University of Aarhus, brought his attention to the beautiful pointer-following strategy. Finding this is left as an exercise in the paper.

Miltersen and Gál originally considered the case where there are n team members and $2n$ lockers, half of them empty; each team member still gets to open up to half of the lockers. This is a more difficult problem. Clearly simple pointer-following will not work as empty lockers do not point anywhere. It is an open question whether the winning probability must tend to zero for large n.

In [5] Navin Goyal and Michael Saks build on Skyum's pointer-following to devise a strategy for Team B in a more general setting, varying both the proportion of empty lockers and the fraction of lockers each team member may open. As the number of players increases, their probability of success for Team B approaches zero less rapidly than conjectured in [4]. And fixing the number of players and fraction of lockers each may open, their probability of winning remains nonzero even as more empty lockers are added.

The problem also appeared in Joe Buhler and Elwyn Berlekamp's puzzle column in the Spring, 2004 issue of *The Emissary* [3], with lockers replaced by ROM locations and colored numbers replaced by signed numbers. Here it is pointed out that the team benefits from the members carefully planning their guessing strategy as well as their locker searching strategy. For example, if there are $2n$ lockers and the longest cycle has length $n + 1$, the team members caught in the $n + 1$ cycle can guess in such a manner that they all guess correctly or all guess incorrectly. The trick is the same as that employed in the hat problem of Todd Ebert [2]. Variations of the hat problem are described in Joe Buhler's article in this column [1] and in Peter Winkler's book [7, p66, p120]. The locker problem will be discussed in a future edition of Winkler's book also.

We thank Joel Spencer for introducing us to the problem, and we thank Ravi Vakil and Michael Kleber for encouraging us to write this note and providing many useful suggestions.

REFERENCES

[1] Joe Buhler, Hat tricks. *Math. Intelligencer* 24 (2002), no. 4, 44–49.

[2] Todd Ebert, http://www.cecs.csulb.edu/~ebert/

[3] *The Emissary*, http://www.msri.org/publications/emissary/

[4] Anna Gál and Peter Bro Miltersen. The Cell Probe Complexity of Succinct Data Structures, *Proceedings of 30th International Colloquium on Automata, Languages and Programming (ICALP)* 2003, 332–344.

[5] Navin Goyal and Michael Saks, A Parallel Search Game, to appear in *Random Structures and Algorithms.*

[6] Richard Stanley, *Enumerative Combinatorics* Vol 1. Cambridge Studies in Advanced Mathematics, 49, Cambridge University Press, Cambridge, 1999

[7] Peter Winkler, *Mathematical Puzzles: A Connoisseur's Collection*, A K Peters, Ltd., Natick, MA, 2004.

AUTHORS

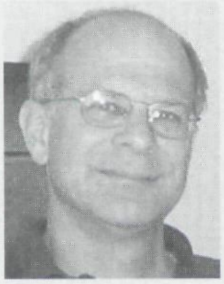

MAX WARSHAUER
Department of Mathematics
Texas State University–San Marcos
San Marcos, TX 78666-4945
e-mail: mw07@academia.swt.edu

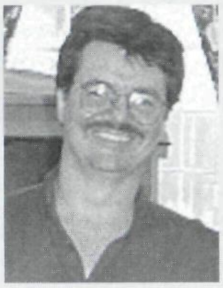

EUGENE CURTIN
Department of Mathematics
Texas State University–San Marcos
San Marcos, TX 78666-4945
e-mail: ec01@txstate.edu

Max Warshauer received his Ph.D. at LSU under Pierre Conner and is now Professor in the Mathematics Department at Texas State, where he founded and directs Mathworks* a center for Mathematics and Mathematics Education, one of five programs in Texas to receive the 2001 Texas Higher Education Star Award for Closing the Gaps. Max was one of 10 individuals in the country to receive the 2001 Presidential Award for Excellence in Science, Mathematics, and Engineering Mentoring. His hobbies include playing chess, go, Ping-Pong, and bicycling. His research interests have included quadratic forms, analysis of algorithms, and mathematics education.

Eugene Curtin, a native of Ireland, received a Ph.D. from Brown University in 1988 under Thomas F. Banchoff. He is currently a professor at Texas State University. His research interests have included differential geometry, ring theory, combinatorics, and graph theory. Eugene has been a faculty member with the Honors Summer Math Camp under Max Warshauer's directorship every summer since 1992. He enjoys mathematical games and puzzles, and likes to play chess, go, and backgammon. He was chess champion of Ireland in 1984 and 1985, and chess champion of Texas in 1991, 1992, and 1998.