References

[1] N. Alon, J. Spencer, *The Probabilistic Method*, John Wiley, New York, 1991

[2] J. Beck, *On* 3-*chromatic hypergraphs*, Discrete Math 24 (1978), pp 127-137

[3] P. Erdős, *On a Combinatorial Problem, I.*, Nordist. Mat. Tidskr. 11 (1963), pp 5-10

[4] P. Erdős,*On a combinatorial problem, II*, Acta. Math. Hungar. 15 (1964), pp 445-447

[5] P. Erdős, L. Lovász, *Problems and results on* 3-*chromatic hypergraphs and some related questions*, in Infinite and Finite Sets, North-Holland, Amsterdam-New York, 1975

[6] P. Erdős and J. Selfridge, *On a combinatorial game*, J. Combinatorial Theory (Ser. A) 14 (1973), pp 298-301

[7] P. Raghavan, *Probabilistic construction of deterministic algorithms: approximating packing integer programs*, J. of Computer and Systems Sciences 37 (1988), pp 130-143

[8] J. Spencer, *Ten Lectures on the Probabilistic Method*, SIAM, Philadelphia, 1987

[9] J. Spencer, *Six Standard Deviations Suffice*, Trans. Amer. Math. Soc. 289 (1985), pp 679-706

[10] J. Spencer, *Balancing Vectors in the Max Norm*, Combinatorica 6 (1986), pp 55-65

Tenure Game, as a finite perfect information game it has a value. Let the initial position consist of a_i faculty i rungs from tenure, as before, and set

$$W = \sum a_i 2^{-i} \tag{34}$$

so that $W = E[T]$ when Carole plays randomly.

Theorem 10 *The value of the generalized Tenure Game is $\lfloor W \rfloor$.*

Proof. Let V, for the moment, be the value and suppose $V > W$. If Carole plays randomly then Paul's expected outcome is W so that his probability of receiving $V > W$ is less than one. In this perfect information Paul can always achieve the value V, a contradiction. Hence $V \leq W$ and, as V is an integer, $V \leq \lfloor W \rfloor$.

For the other side we use the following: Let $x_1 \geq \ldots \geq x_l$ be negative powers of two with sum at least K, where K is an integer. Then there is a partition of the x_i into two groups so that each group sums to at least $K/2$. Let's assume this simple corollary of the Splitting Lemma. Set $K = \lfloor W \rfloor$. Now, as with Theorem 9, Paul always creates a list L so that $E[T^1] \geq K/2$ and $E[T^2] \geq K/2$ and at the end of the game $E[T] \geq K$ so that Paul has pushed K faculty into tenured positions.

Our final example is a reversal of the generalized tenure game. The rules remain the same except that now the value of the game to *Carole* is the number of faculty receiving tenure. Here the Chair Paul is the bad guy trying to prevent faculty receiving tenure. Call this the Good Dean game.

Theorem 11 *The value of the Good Dean game is $\lceil W \rceil$.*

Proof. Let V, for the moment, be the value and suppose $V < W$. If Carole plays randomly then Carole's expected outcome is W so that the probability of Carole receiving $V < W$ is less than one. In this perfect information Paul can always hold Carole to at most the value V, a contradiction. Hence $V \geq W$ and, as V is an integer, $V \geq \lceil W \rceil$.

For the other side we use the following: Let $x_1 \geq \ldots \geq x_l$ be negative powers of two with sum at most K, where K is an integer. Then there is a partition of the x_i into two groups so that each group sums to at most $K/2$. Let's assume this simple corollary of the Splitting Lemma. Set $K = \lceil W \rceil$. Now, as with Theorem 9, Paul always creates a list L so that $E[T^1] \leq K/2$ and $E[T^2] \leq K/2$ and at the end of the game $E[T] \leq K$ so that Paul has held the number of faculty in tenured positions to at most K.

groups so that each group sums to at precisely one half.
Proof. We place the x_i into groups largest first, always placing x_i into the group with the currently smaller sum. Let us say we are stuck at l if after placing $x_1, \ldots, x_l$ the difference of the sums of the groups (in absolute value) is greater than the sum $x_{l+1}+\ldots+x_r$ of the as yet unplaced x's. We show by induction on l, $0 \leq l \leq r$, that we are never stuck. We are trivially not stuck at $l=0$, assume by induction that we are not stuck at $l-1$. Case 1: the two groups currently have different sums. As all $x_1, \ldots, x_{l-1}$ are multiples of x_l the difference of the sums of the groups must be a multiple of x_l. Hence the diffence is at least x_l and so placing x_l in the smaller group cannot make us stuck. Case 2: the two groups currently have the same sum. This sum, as in Case 1, must be of the form Ax_l, A integral. Thus $x_1+\ldots+x_l$ is of the form $(2A+1)x_l$ and hence

$$x_{l+1}+\ldots+x_r = 1-(2A+1)x_l \geq x_l \tag{32}$$

so that after placing x_l in either group we are not stuck. Hence we will not be stuck at $l=r$ which means that after placement of all $x_1, \ldots, x_l$ the sums are precisely the same. □
Corollary. Let $x_1 \geq \ldots \geq x_l$ be negative powers of two with sum at least one. Then there is a partition of the x_i into two groups so that each group sums to at least one half.
Proof. If $x_1+\ldots+x_l > 1$ then, since it is a multiple of x_l, $x_1+\ldots+x_{l-1} \geq 1$. Remove $x_l, x_{l-1}, \ldots$ until $x_1+\ldots+x_r = 1$ and apply the Splitting Lemma. □

Theorem 9 *If*

$$\sum a_k 2^{-k} \geq 1 \tag{33}$$

then Paul wins the Tenure Game.

Proof. Initially $E[T] \geq 1$. From the Lemma Paul may create a list L so that $E[T^1] \geq 1$ and $E[T^2] \geq 1$. (Note that $E[T^1]$ is defined after Carole plays Option One and so is double the sum of the original weights of the faculty in list L.) Regardless of what Carole does $E[T] \geq 1$ at the end of the round. At the end of the game $E[T] \geq 1$ and thus someone has received tenure and Paul has won. □

Lets generalize a bit. The rules remain the same except that now the value of the game to Paul is the *number* of faculty receiving tenure. That is, Chair Paul wants to maximize the number of faculty receiving tenure and Dean Carole wants to minimize that number. Call this the generalized as those f which are k rungs from tenure each have $E[I_f] = 2^{-k}$. Note that Carole wins if and only if $T = 0$. Our assumption 28 may be restated that $E[T] < 1$ and hence

$$\Pr[\text{Carole wins}] = \Pr[T = 0] > 0 \tag{30}$$

Now comes the slick part. The Tenure Game is a finite perfect information game with no draws so that either Paul or Carole has a perfect strategy. Had Paul had a perfect strategy then by playing it the probability of Carole winning would be zero, which is not the case. Hence, Carole *must* have a winning strategy! □

The above proof is a nice example of the probabilistic method, the use of probabilistic analysis to prove a deterministic result. As often the case with the probabilistic method it leaves open the question of actually finding the desired object – in this case Carole's strategy.

Proof 2. Define the weight of a position as the expected number $E[T]$ of faculty receiving tenure if Carole plays randomly. Explicitly, with a_k faculty k rungs from tenure the weight is $\sum a_k 2^{-k}$. Now Paul presents a list L to Carole. Let T^1 be the number of faculty receiving tenure if Carole now plays Option One and then plays randomly in all succeeding rounds. Let T^2 be the same with Carole first playing Option Two. Carole's strategy is to pick Option One if $E[T^1] < E[T^2]$, otherwise to pick Option Two. (Suppose Option One leaves b_k players k rungs from tenure after its application while Option Two leave c_k players k rungs from tenure. Then Carole simply checks if $\sum b_k 2^{-k} < \sum c_k 2^{-k}$ and hence this is a very efficient strategy.) The key point here is that

$$E[T] = \frac{1}{2}\left(E[T^1] + E[T^2]\right) \tag{31}$$

since playing randomly throughout is the average of playing Option One and then randomly and playing Option Two and then randomly. As $E[T] < 1$ either $E[T^1] < 1$ or $E[T^2] < 1$ and employing this strategy Carole insures that $E[T] < 1$ at the end of the round. Continuing this at the end of the game $E[T] < 1$. But at the end of the game $E[T]$ is simply the number of faculty who have received tenure. An integer less than one must be zero so Carole has won. □

The Tenure Game has the nice property that when the condition for Carole winning does not hold Paul can use this same weight function to give a winning strategy for himself. We need in this case an amusing lemma.

Splitting Lemma. Let $x_1 \geq x_2 \geq \ldots \geq x_r$ all be negative powers of two with sum $x_1 + \ldots + x_r = 1$. Then there exists a partition of the x_i into two tenure – Carole wins if no faculty member receives tenure. Each year (or round if you will) Chair Paul creates a promotion list L, a subset of the faculty[2], and gives it to Dean Carole who has two options. Option One: Carole may promote all faculty on list L one rung and simultaneously fire all other faculty. Option Two: Carole may promote all faculty *not* on list L one rung and simultaneously fire all faculty on list L. With the example above, suppose $L = \{WALTER, JARIK\}$. If Carole applies Option One WALTER receives tenure and Paul has won. So Carole would apply Option Two: WALTER and JARIK would disappear, RON and BELA would become level two Assistant Professors and INGO and HANS would become level one Assistant Professors. The next year Paul presents another list L and Carole picks one of the two options. The Tenure Game represents an extreme form of "publish or perish", within four years all faculty will either have been promoted to tenure or fired. With perfect play on both sides, who wins the Tenure Game?

Naturally we shall consider a general opening position, let us suppose that there are a_k faculty that are k rungs from tenure and that k can be arbitrarily large, though bounded.

Theorem 8 *If*

$$\sum a_k 2^{-k} < 1 \tag{28}$$

then Carole wins.

Proof 1. Let us imagine that Carole plays *randomly*, i.e., each round after Paul has determined the promotion list L Carole flips a fair coin to decide whether to use Option 1 or Option 2. Fix some deterministic strategy for Paul. Now each faculty has a probability of reaching tenure – for the example above RON has probability $\frac{1}{8} = 2^{-3}$ of receiving tenure since for the next three years Carole must select the Option that promotes, rather than fires, RON. Note critically that this probability is 2^{-3} regardless of Paul's strategy; when Paul puts RON in L Carole must choose Option One while when Paul leaves RON out of L Carole must choose Option Two but each occurs with probability $\frac{1}{2}$. Let T be the number of faculty receiving tenure so that T is a random variable. For each faculty member f let I_f be the indicator random variable for f receiving tenure so that $T = \sum I_f$. Then by Linearity of Expectation

$$E[T] = \sum E[I_f] = \sum a_k 2^{-k} \tag{29}$$

[2]The faculty are only pawns in this game!

formulation) but simultaneously create $S_{n+i} = \overline{S_i}$, $1 \le i \le n$. Stop after n rounds, as this can only help Paul. Some

$$
\begin{aligned}
|\chi^*(S_i)| &= |\chi(S_i) - \chi(S_{n+i})| \\
&\ge VAL^*(n) > c\sqrt{n}\sqrt{\ln n}
\end{aligned} \tag{25}
$$

and so for some $1 \le i \le 2n$

$$
|\chi(S_i)| \ge \frac{1}{2} c\sqrt{n}\sqrt{\ln n} \tag{26}
$$

Putting the results together we have the best result (originally given with a different proof, in [10]) up to constant factor:

Theorem 7 $VAL(n) = \Theta(\sqrt{n}\sqrt{\ln n})$

In 1985 this author [9] improved Theorem 3 and showed that given any family $\mathcal{F}$ consisting of n sets on n points there was a coloring χ with

$$
disc(\mathcal{F}, \chi) < 6\sqrt{n} \tag{27}
$$

(The 6 was a constant which was improvable by more precise calculations.) Unlike the situation with the weaker $O(\sqrt{n \ln n})$ bound, there is no polynomial time algorithm known that will find such a coloring χ and it seems quite possible to this author that no such algorithm exists. The results on $VAL(n)$, while certainly not contradicting the possibility of such an algorithm, argue that such an algorithm would have to be "global", there could be no such algorithm that considered and colored the points one at a time.

3 The Tenure Game

$HANS$				
$INGO$	RON			
$JARIK$	$BELA$		$WALTER$	
_ _ _ _ _	_ _ _ _	_ _ _	_ _ _ _	_ _ _ _
$PostD$	$AP1$	$AP2$	$Assoc$	$Tenure$

The tenure game is a perfect information game between two players, Paul - chairman of the department - and Carole - dean of the school. An initial position is given in which various faculty (HANS , INGO, etc.) are at various pre-tenured positions. Paul will win if some faculty member receives

$(a_{1x}, \ldots, a_{nx})$ the values w^R, w^B satisfy

$$w^R - w^B = \sum_{i=1}^{n} a_{ix}(w_i^R - w_i^B) \tag{19}$$

Carole will select a_{ix} to minimize the absolute value of this sum. While computing the minimum may be difficult a simple greedy algorithm allows Carole to select the a_{ix} so that

$$|w^R - w^B| \leq \max_i |w_i^R - w_i^B| \tag{20}$$

Let $y = n - 1 - x$, be the number of rounds remaining after the x-th round. How much can the choice of ϵ_x affect the weight of i? Some thought shows that $w_i^R - w_i^B$ can be at most the probability that the random $\epsilon_{x+1} + \ldots + \epsilon_n \sim Z_y$ is a particular value. This happens with probability at most $\binom{y}{y/2} 2^{-y} = O(y^{-1/2})$, with the understanding that when $y = 0$ this is at most one. That is:

$$|w^R - w^B| = O(y^{-1/2}) \tag{21}$$

Let w^O, w^N be the weights before and after the x-th round, with Carole using this strategy. Then regardless of what Paul does

$$w^N \geq \min[w^R, w^B] = w^O - \left|\frac{w^r - w^B}{2}\right| = w^O - O(y^{-1/2}) \tag{22}$$

Let w^I and w^F be the initial and final weights respectively. Then

$$w^F \geq w^I - \sum_{y=0}^{n-1} O(y^{-1/2}) \geq w^I - O(\sqrt{n}) \tag{23}$$

Now suppose n, α are such that $w^I > c\sqrt{n}$ with c the constant of the O-term above. Then with Carole playing this strategy $w^F > 0$. But w^F is simply the number of i with $|L_i| > \alpha$ so there would be at least one such i and therefore Carole would win the game. Here

$$w^I = n \Pr[|Z_n| > \alpha] \tag{24}$$

with Z_n given by 9. Again Large Deviation results can be used to show that with $n = c\sqrt{n}\sqrt{\ln n}$, c a small constant, $w^I \gg n^{1/2}$ and so Carole wins. $\square$

From this we may easily derive a lower bound on $VAL(2n)$. Let Carole create sets S_i, $1 \leq i \leq n$ according to the $VAL^*(n)$ game (in the set on the x-th round Carole decides the profile of x and then Paul decides $\chi(x) = \pm 1$. Now, however, we define

$$\chi^*(S_i) = \chi(S_i) - \chi(\overline{S_i}) \tag{16}$$

and let the payoff to Carole be the maximal $|\chi^*(S_i)|$ over $1 \leq i \leq n$. Let $VAL^*(n)$ denote the value to Carole of this new game. It will also be convenient to describe this game geometrically. On the x-th round Carole selects a vector $v_x = (a_{1x}, \ldots, a_{nx})$ with all $a_{ix} = \pm 1$. Paul then selects $\epsilon_x = \pm 1$. Set

$$L_i = \epsilon_1 a_{i1} + \ldots + \epsilon_n a_{in} \tag{17}$$

so that

$$\epsilon_1 v_1 + \ldots + \epsilon_n v_n = (L_1, \ldots, L_n) \tag{18}$$

The payoff to Carole is then the maximal absolute value of the L_i or, in geometric terms, the L^∞ norm of $\epsilon_1 v_1 + \ldots + \epsilon_n v_n$. The equivalence between the descriptions is seen by setting $\epsilon_x = \chi(x)$ and corresponding $a_{ix} = +1$ with $x \in S_i$ and $a_{ix} = -1$ with $x \notin S_i$. (Perhaps the game is best thought of dynamically. There is a position vector $P \in R^n$, initially set to 0. There are n rounds. On each round Carole selects a vector $v \in R^n$ with all coordinates ± 1 and then Paul resets P to be $P \pm v$, his choice of sign. The payoff to Carole is the L^∞ norm of P at the end of the game.)

Theorem 6 $VAL^*(n) = \Omega(\sqrt{n}\sqrt{\ln n})$

Proof. At an intermediate position define the weight of i to be the probability that at the end of the game $|L_i| > \alpha$ if the remaining ϵ_x are randomly chosen ± 1. (Note the a_{ix} are then immaterial since $\epsilon_x a_{ix} = \pm 1$ randomly.) The weight of a position is then the sum over $1 \leq i \leq n$ of the weight of i. (If the weight of the initial position is less than one then one can argue as before that Paul will win but now we want Carole to win.) Let w^O be the weight (we'll leave off the asterisks) at the beginning of the x-th round. A move by Carole would generate values w^R, w^B of new weights should Paul select $\epsilon_x = +1$ or -1 respectively. As before $w^O = \frac{1}{2}(w^R + w^B)$. Carole's strategy will be to make that move so that w^R and w^B are as close together as possible, i.e., to minimize $|w^R - w^B|$.

Let w_i^R, w_i^B be the weights of i *if* Carole selects $a_{ix} = +1$ and Paul selects $\epsilon_x = +1$ or -1 respectively. If, instead, Carole were to select $a_{ix} = -1$ the roles of w_i^R and w_i^B would be reversed. Summing over i, when Carole selects smallgame is entirely inside Paul's mind!) We claim

$$\frac{1}{2}(w^R + w^B) \leq w^O \tag{13}$$

Once again, by additivity, it suffices to show this where the w's are the contributions by a single S_j. If smallgame (beginning with x) is restricted by requiring Carole to pick $x \in S_j$ then Carole's optimal probability of winning is $\frac{1}{2}(w^R + w^B)$ as after the coin flip for $\chi(x)$ Carole will play optimally. Removing that restriction gives the actual game starting at x and Carole's optimal probability w^O when she has no restrictions cannot be smaller than when she has restriction.

As with 5, this implies $w^N \leq w^O$.

As before, let w^I, w^F denote the initial and final weights respectively. Assume $w^I < 1$. Then $w^F \leq w^I < 1$. In the final position the game is over so that the smallgame has no rounds to it and w^F is simply the number of S_j with $|\chi(S_j)| > \alpha$. This nonnegative integer is less than one and therefore zero and therefore Paul would have won.

The above paragraphs are very general, now we must estimate w^I in terms of n and α. By symmetry $w^I = np$ where p is Carole's optimal probability of winning smallgame from the initial position. Determination of p is known to probabilists as a stopping rule problem and can be expressed more colorfully in gambling language. Carole begins with \$0. but unlimited credit. There are n rounds and on each round she can either pass ($x \notin S_j$) or bet \$1. at even odds. Her Dostoyevskian goal is to end with either a win or a loss greater than α. It is intuitively clear (and rigorously provable) that her best strategy is to always bet until she reaches her goal and then pass forevermore. Let $X_i = \pm 1$ be uniform and independent and set

$$M_n = \max_{1 \leq i \leq n} |X_1 + \ldots + X_i| \tag{14}$$

Carole wins if and only if $M_n > \alpha$ so that

$$w^I = pn = n \Pr[M_n > \alpha] \tag{15}$$

We have actually shown that if $\Pr[M_n > \alpha] < n^{-1}$ then Paul wins. From elementary probability (the reflection principle) one can show $\Pr[M_n > \alpha] < 4\Pr[Z_n > \alpha]$ with Z_n given 9. From 10 the condition $w^I < 1$ then holds for $\alpha = \sqrt(2n)\sqrt(\ln(8n))$. $\square$

Now let us "reverse" the game analysis to give a lower bound on $VAL(n)$. It is technically convenient to consider a slightly different game. As before

before $w^O = \frac{1}{2}(w^R + w^B)$ and so the new weight w^N has $w^N \le w^O$. At the end of this procedure the final position has weight $w^F \le w^I < 1$. But w^F is simply the number of S_j with $|\chi(S_j)| > \alpha$. A nonnegative integer less than one must be zero so that χ is the desired coloring. □

This approach to derandomization of an algorithm via a weight function has been more fully developed in Spencer [8] and Raghavan [7]. Now we resurrect Carole and make a game version of Theorem 4. At the end of the game there will be a family $S_1, \ldots, S_n$ of subsets of $\{1, \ldots, n\}$ but at the beginning these have not been determined. On the j-th round Carole gives the *profile* of point j - i.e., she says if $j \in S_i$ for each $1 \le i \le n$ - and then Paul determines $\chi(j)$. At the end of the game $\mathcal{F}$ and χ have been determined and the payoff, to Carole, is $disc(\mathcal{F}, \chi)$. Carole need not have a particular family of S_i in mind at the start of the game but rather gives the best profile in each round given the history. This is a perfect information game and so it has a value, call it $VAL(n)$, to Carole.

Theorem 5 $VAL(n) = O(\sqrt{n}\sqrt{\ln n})$

Proof. We fix $\alpha = c\sqrt{n}\sqrt{\ln n}$ for an appropriately large c and give a strategy for Paul that assures that at the end of the game all $|\chi(S_j)| < \alpha$. Again there is a weight function but this time there is a new twist. Consider an intermediate position and a particular S_j which has some $\chi(i) = \pm 1$ and others undetermined.

Consider the following solitaire game by Carole, lets call it *smallgame*. For each undetermined i, sequentially, Carole decides if $i \in S_j$ and then a fair coin is flipped to determine $\chi(i) = \pm 1$. Carole wins smallgame if at the end of the game $|\chi(S_j)| > \alpha$. (By way of illustration, if the sum of the already determined $\chi(i)$ with $i \in S_j$ is bigger than α at the start then Carole can say that no new i are in S_j and so she will win with probability one.) Define the weight of S_j to be the probability that Carole will win smallgame, assuming she plays an optimal strategy. As before, the weight of a position is the sum of the weights of the S_j.

At some intermediate position call the weight, as defined above, w^O. In the next round Carole decides the profile of, say, x and then Paul determines $\chi(x)$. Paul calculates w^R and w^B, what the new weights would be should he select $\chi(x) = +1$ or $\chi(x) = -1$ respectively. Paul's strategy is to select $\chi(x)$ so as to minimize the new weight, which we call w^N. (Let us be clear that we are making no assumptions about what strategy Carole really uses,

Theorem 3 *There is a χ for which*

$$disc(\mathcal{F},\chi) \leq \alpha = (2n\ln(2n))^{1/2} \tag{8}$$

Proof. We require a basic fact about Large Deviations. Let $X_1,\ldots,X_a$ be independent random variables with $\Pr[X_i = +1] = \Pr[X_i = -1] = \frac{1}{2}$ and set

$$Z_a = X_1 + \ldots + X_a \tag{9}$$

(This is usually denoted S_a but here we use S for sets.) Then (see, e.g., the appendix of [1]) for any positive β

$$\Pr[|Z_a| > \beta] < 2e^{-\beta^2/2a} \tag{10}$$

Now let χ be random. For each $1 \leq j \leq n$ the random variable $\chi(S_j)$ has distribution Z_a with $a = |S_j|$. Letting A_j be the event $|\chi(S_j)| > \alpha$

$$\Pr[A_j] < 2e^{-\alpha^2/2a} \leq 2e^{-\alpha^2/2n} = \frac{1}{n} \tag{11}$$

by the choice of α. Letting A be the disjunction of the A_j,

$$\Pr[A] \leq \sum \Pr[A_j] < n\frac{1}{n} = 1 \tag{12}$$

and so the event $\overline{A}$ has positive probability. Thus there is a point in the probability space, a particular χ, for which A fails - i.e., for which all $|\chi(S_j)| \leq \alpha$. □

The weight function approach can be easily modified to handle this result

Theorem 4 *There is a polynomial time algorithm to find the χ guaranteed by Theorem 3.*

Proof. The values $\chi(i)$, $1 \leq i \leq n$, will be found by Paul sequentially. At an intermediate stage some $\chi(i) = +1$, some $\chi(i) = -1$ and some $\chi(i)$ are undetermined. The weight of a set S_j is defined as the probability that $|\chi(S_j)| > \alpha$ if the undetermined $\chi(i)$ are set equal ± 1 randomly. Observe that this weight, being the sum of appropriate binomial coefficients, is rapidly calculatable. The weight of a position is the sum of the weights of the S_j. Initially each S_j has weight less than $\frac{1}{n}$ so that the initial position has weight $w^I < n\frac{1}{n} = 1$. Consider a position with x the next point to be colored. As before let w^O be the weight of the current position and let w^R, w^B be the new weights if $\chi(x)$ is set equal to $+1$ or -1 respectively. As some Blue points, no Red points, and s uncolored points its weight is again 2^{-s}. If, as in the beginning, it has n uncolored points it has weight 2^{1-n}. The weight of a position is the sum of the weights of the S_j. The initial position thus has weight $m2^{1-n} < 1$ and if the final position contains any monochromatic sets it will have weight at least one. Paul colors a new $x \in \Omega$ Red or Blue so as the minimize the weight of the new position. Let w^O be the weight before coloring x and w^R, w^B the new weights if x is colored Red or Blue respectively. We claim

$$w^O = \frac{1}{2}(w^R + w^B) \tag{4}$$

By addititivity it suffices to look at the weight of any particular S_j. The contribution to w^R is the probability S_j becomes monochromatic if x is colored Red and the remaining points randomly. The contribution to w^B is the probability S_j becomes monochromatic if x is colored Blue and the remaining points randomly. Then their average is the probability S_j becomes monochromatic if x is colored randomly and then the remaining points randomly, precisely the contribution to w^O. Let w^N be the weight after coloring x. Paul's strategy assures

$$w^N = \min(w^R, w^B) \le \frac{1}{2}(w^R + w^O) = w^O \tag{5}$$

Let w^I, w^F denote the initial and final weights respectively. As the weight never increases $w^F \le w^I < 1$ and so the final coloring has no monochromatic sets. Ignoring points in no sets we have m sets on at most mn points. This algorithm takes polynomial (in m, n) time to find the desired coloring. Observe that an exhaustive search could take time 2^m or more.

2 Discrepancy

Now let us restrict $\mathcal{F}$ to be a family of n sets $S_1, \ldots, S_n$ of undetermined sizes, all subsets of an n-set $\Omega = \{1, \ldots, n\}$. Here the object will be to color so that the sets S_j are evenly balanced. For any $\chi : \Omega \to \{-1, +1\}$ define

$$\chi(S_j) = \sum_{x \in S_j} \chi(x) \tag{6}$$

and define the discrepancy $disc(\mathcal{F}, \chi)$ of $\mathcal{F}$ with respect to χ by

$$disc(\mathcal{F}, \chi) = \min_{1 \le j \le n} |\chi(S_j)| \tag{7}$$

The following bound is due to Erdős.

by Paul has weight $2^{-0} = 1$. The weight of a position is defined as the sum of the weights of the S_j. Initially all sets have weight 2^{-n} and so the initial position has weight $m2^{-n} < 1$.

Carole's strategy is always to select that previously unselected $x \in \Omega$ so that the weight of the new position is minimized. Consider a single round in which Carole selects x and then Paul selects y. The selection of x decreases the position weight by the sum of all the weights of the sets containing x. The selection of y increases the position weight by the sum of all the weights of the sets containing y but not x. (If a set had had s unselected points, with weight 2^{-s}, the selection of y by Paul increases its weight by 2^{-s} to $2^{-(s-1)}$.) With Carole going first and optimizing the decrease by selecting x is no less than the increase by Paul's selecting y; thus the weight at the end of the round is not greater than when the round began. The initial position has weight less than one, hence the final position has weight less than one. But had Paul selected an entire S_j the weight would be at least one. Thus Paul has lost and therefore Carole has won. $\square$

Erdős and Selfridge also show that the condition $m < 2^n$ is best possible. For let $\Omega = \{1, \ldots, 2n\}$ and consider the "Chinese Menu" $\mathcal{F}$ consisting of all n-sets $S \subset \Omega$ which have, for all $1 \leq t \leq n$, precisely one of the pair $2t-1, 2t$. This family has 2^n sets and Paul wins by the natural pairing strategy. The situation with Theorem 1 is less clear. Erdős has defined $m(n)$ to be the minimal size of a family $\mathcal{F}$ of n-sets having the property that every coloring leaves some $S_j \in \mathcal{F}$ monochromatic. Theorem 1 then gives $m(n) \geq 2^{n-1}$ but this is not best possible. The best known asymptotic bounds on $m(n)$ are

$$c_1 2^n n^{1/3} < m(n) < c_2 2^n n^2 \tag{3}$$

due to Jozsef Beck [2] and Erdős [4] respectively. These bounds, while leaving a fascinating and important problem, will not further concern us here.

The weight function approach of Erdős and Selfridge can be adapted to give an algorithm implementation of Theorem 1. Again let $\mathcal{F}$ consist of m sets $S_j \subset \Omega$, each of size n, with $m < 2^{n-1}$. Ω is ordered in some way. Paul (Carole disappears for a while) wants to find the coloring guaranteed by Erdős. Paul colors the points sequentially. At an intermediate position there are Red points, Blue points and uncolored points. Define the weight of S_j at this position to be the probability that it will be monochromatic if its remaining uncolored points are randomly colored Red and Blue. That is, if S_j already has a Red and a Blue point its weight is zero. If it has some Red points, no Blue points, and s uncolored points its weight is 2^{-s}. If it has if all the $\chi(x)$, $x \in S_j$ are the same. In 1963 Erdős [3] gave the following result.

Theorem 1 *Let $\mathcal{F} = \{S_1, \ldots, S_m\}$ with all $|S_j| = n$ and with $m < 2^{n-1}$. Then there exists a coloring χ so that no S_i is monochromatic.*

Proof. Let χ be a *random* coloring of Ω. That is, for each $x \in \Omega$,

$$\Pr[\chi(x) = +1] = \Pr[\chi(x) = -1] = \frac{1}{2} \tag{1}$$

and the values $\chi(x)$ are mutually independent. (We "flip a fair coin" to determine each $\chi(x)$.) To every set S_j we correspond the *event* A_j that S_j is monochromatic. Clearly $\Pr[A_j] = 2^{1-n}$. Let A be the event that some S_j is monochromatic, so that $A = \vee A_j$. Then

$$\begin{aligned} \Pr[A] = \Pr[\vee_{j=1}^{m} A_j] &\leq \sum_{j=1}^{m} \Pr[A_j] \\ &= m2^{1-n} < 1 \end{aligned} \tag{2}$$

given the assumption on m. The event $\overline{A}$ then has nonzero probability. Hence there is a point in the probability space - i.e., a coloring χ - for which no S_j is monochromatic. □

The natural, if vague, question arises: where is the coloring? Before approaching this from an algorithmic viewpoint we turn to a seminal paper [6] of Erdős and John Selfridge in 1973. They considered a game between two players, who we will call Paul and Carole. A family $\mathcal{F} = \{S_1, \ldots, S_m\}$ on a set Ω is given, visible to both players. On each round Carole selects an $x \in \Omega$ and then Paul selects a $y \in \Omega$. Once selected a point may not be selected again. The game ends when all of Ω has been selected. Paul wins if there is an S_j for which he has selected *all* $x \in S_j$. Carole wins if Paul doesn't win.

Theorem 2 *Assume all $|S_j| = n$ and $m < 2^n$. Then Carole wins.*

Proof. Consider an intermediate position in the game, when points have either been selected by Paul, been selected by Carole, or remain unselected. Define the *weight* of S_j in that position to be zero if any points have been selected by Carole (i.e., the set has been "killed") and, critically, to be 2^{-s} if no points have been selected by Carole and there remain s unselected points. This includes the case $s = 0$ so that a set which has been entirely selected

From Erdős to Algorithms

Joel Spencer[1]

How can one explain the revolutionary rise of the Discrete over the past half century. Conditions had to be ripe. Mathematical proof may be absolute but the directions of mathematical thought are buffetted by the winds of social change. We were hit by a full scale hurricane: The Computer. It has changed *everything*. The way we do mathematics has changed, from the methods of printing papers to the obtaining of approximate solutions to partial differential equations. The story for this author is the change in the paradigms of Mathematics itself. Algorithms have come to center stage. Hilbert has lost. Existence is no longer enough. Even the recursiveness of the first half of our century is no longer enough. We now want to avoid *intractibility*, we want a *polynomial time* algorithm for constructing our objects.

Social conditions do not, by themselves, lead to change. There must be leaders to catch the winds. Paul Erdős has done this for Discrete Mathematics. Through his travelling, his discussions, his theorems and, perhaps most importantly, his conjectures large areas of Discrete Mathematics have been developed. Here we concentrate on just one of the many areas he has developed: The Probabilistic Method. There is an irony here. Erdős himself has never programmed a computer and rarely speaks of algorithmic questions. In his own papers to a large extent he holds a Hilbertian philosophy, he proves the existence of the desired coloring, tournament, design, graph or whatever and then he moves on to the next problem. The methods he has developed have been redesigned to fit the Algorithmic paradigms. The Probabilistic Method, and more generally the use of randomness in algorithms, has proven to be a central idea in Theoretical Computer Science.

But enough chit chat. Lets follow the Erdős ideal and look at specific problems with specific solutions.

1 Monochromatic Sets

Our general format will consist of a universal set Ω and a family $\mathcal{F} = \{S_1, \ldots, S_m\}$ of subsets of Ω. A coloring is a map $\chi : \Omega \to \{-1, +1\}$. (We'll often identify $+1$ with Red and -1 with Blue. *Throughout this paper all colorings shall be with* two colors.) A set S_j is monochromatic (under χ)

[1] Adapted from Lectures at the SODA conference, Orlando, FL; SIAM conference, Vancouver, BC Canada; and the Bielefeld meeting, all 1992