5 Recoloring

Suppose that a random coloring leaves a set of blemishes. Here we apply a random recoloring to the blemishes to remove them. If the recoloring is too weak then not all the blemishes are removed. If the recoloring is too strong then new blemishes are created. The recoloring is given a parameter p and these too possibilities are decreasing and increasing functions of p. Calculus then points us to the optimal p.

We use the notation of Chapter. 1, §3 on Property B: $m(n) > m$ means that given any n-uniform hypergraph $H = (V, E)$ with m edges there exists a 2-coloring of V so that no edge is monochromatic. J. Beck [1978] improved Erdős's 1963 bound to $m(n) = \Omega(2^n n^{1/3})$. Building on his methods Srinivasan and Radhakrishnan [1998] proved $m(n) = \Omega(2^n (n/\ln n)^{1/2})$ and it is that proof we shall give. While this proof is neither long nor technically complex it has a number subtle and beautiful steps and it is not surprising that it took thirty five years to find it. That said, the upper and lower bounds on $m(n)$ remain quite far apart!

Theorem 5.1. If there exists $p \in [0, 1]$ with

$$k(1-p)^n + k^2 p < 1$$

then $m(n) > 2^{n-1}k$.

Corollary 5.2. $m(n) = \Omega(2^n (n/\ln n)^{1/2})$

Proof. Bound $1 - p \le e^{-p}$. The function $ke^{-pn} + k^2 p$ is maximized at $p = \ln(n/k)/n$. Substituting back in, if

$$\frac{k^2}{n}\left[1 + \ln(n/k)\right] < 1$$

then the condition of Theorem 5.1 holds. This inequality is true when $k = c(n/\ln n)^{1/2}$ for any $c < \sqrt{2}$ with n sufficiently large. □

The condition of Theorem 5.1 is somewhat typical, one wants the total failure probability to be less than one and there are two types of failure. Oftentimes one finds reasonable bounds by requiring the stronger condition that each failure type has probability less than one half. Here $k^2 p \le \frac{1}{2}$ gives $p \le \frac{1}{2}k^{-2}$. Plugging the maximal possible p into the second inequality $k(1-p)^n \le \frac{1}{2}$ gives $k^2 \ln(2k) \le n$. This again holds when $k = c(n/\ln n)^{1/2}$ though now we have the weaker condition $c < 1/\sqrt{2}$. We recommend this rougher approach as a first attempt at a problem, when the approximate range of the parameters is still in doubt. The refinements of calculus can be

placed in the published work!
Proof of Theorem 5.1 Fix $H = (V, E)$ with $m = 2^{n-1}k$ edges and p satisfying the condition. We describe a randomized algorithm that yields a coloring of V. It is best to preprocess the randomness: Each $v \in V$ flips a first coin, which comes up heads with probability $\frac{1}{2}$ and a second coin, which comes up heads (representing potential recoloration) with probability p. In addition (and importantly), the vertices of V are ordered randomly.
Step 1. Color each $v \in V$ red if its first coin was heads, otherwise blue. Call this the first coloring. Let D (for dangerous) denote the set of $v \in V$ that lie in some (possibly many) monochromatic $e \in E$.
Step 2. Consider the elements of D sequentially in the (random) order of V. When d is being considered call it still dangerous if there is a (possibly many) $e \in H$ containing d that was monochromatic in the first coloring and for which no vertices have yet changed color. If d is not still dangerous then do nothing. But if it is still dangerous then check its second coin. If it is heads then change the color of d, otherwise do nothing. We call the coloring at the time of termination the final coloring.

We say the algorithm fails if some $e \in H$ is monochromatic in the final coloring. We shall bound the failure probability by $k(1-p)^n + k^2 p$. The assumption of Theorem 5.1 then assures us that with positive probability the algorithm succeeds. This, by our usual magic, means that there is some running of the algorithm which yields a final coloring with no monochromatic e, that there exists a 2-coloring of V with no monochromatic edge. For convenience, we bound the probability that some $e \in H$ is red in the final coloring, the failure probability for the algorithm is at most twice that.

An $e \in E$ can be red in the final coloring in two ways. Either e was red in the first coloring and remained red through to the final coloring or e was not red in the first coloring but was red in the final coloring. [The structure of the algorithm assures us that points cannot change color more than once.] Let A_e be the first event and C_e the second. Then

$$\Pr[A_e] = 2^{-n}(1-p)^n$$

The first factor is the probability e is red in the first coloring, that all first coins of e came up heads. The second factor is the probability that all second coins came up tails. If they all did then no $v \in e$ would be recolored in Step 2. Inversely, if any second coins of $v \in e$ came up heads there would be a *first* v (in the ordering) that came up heads. When it did v was still dangerous as e was still monochromatic and so v does look at its second coin and change its color. We have

$$2\sum_{e\in H}\Pr[A_e] = k(1-p)^n$$

giving the first addend of our failure probability.

In Beck's 1978 proof, given in our first edition, there was no notion of still dangerous - every $d \in D$ changed its color if and only if its second coin was heads. The values $\Pr[A_e] = 2^{-n}(1-p)^n$ are the same in both arguments. Beck had bounded $\Pr[C_e] \le k^2 pe^{pn}$. The new argument avoids excessive recoloration and leads to a better bound on $\Pr[C_e]$. We turn to the ingenious bounding of $\Pr[C_e]$.

For distinct $e, f \in E$ we say e *blames* f if

- e, f overlap in precisely one element. Call it v.
- In the first coloring f was blue and in the final coloring e was red.
- In Step 2 v was the *last* vertex of e that changed color from blue to red.
- When v changed its color f was still entirely blue.

Suppose C_e holds. Some points of e changed color from blue to red so there is a *last* point v that did so. But why did v flip its coin? It must have been still dangerous. That is, v must be in some (perhaps many) set f that was blue in the first coloring and was still blue when v was considered. Can e, f overlap in another vertex v'. No! For such a v' would necessarily have been blue in the first coloring (as $v' \in f$) and red in the final coloring (as $v' \in e$), but then v' changed color before v. Hence f was no longer entirely blue when v was considered and so e could not blame f. Therefore, when C_e holds e blames some f. Let B_{ef} be the event that e blames f. Then $\sum_e \Pr[C_e] \le \sum_{e\neq f}\Pr[B_{ef}]$. As there are less than $(2^{n-1}k)^2$ pairs $e \neq f$ it now suffices to bound $\Pr[B_{ef}] \le 2^{1-2n}p$.

Let e, f with $e \cap f = \{v\}$ (otherwise B_{ef} cannot occur) be fixed. The random ordering of V induces a random ordering σ of $e \cup f$. Let $i = i(\sigma)$ denote the number of $v' \in e$ coming before v in the ordering and let $j = j(\sigma)$ denote the number of $v' \in f$ coming before v in the ordering. Fixing σ we claim

$$\Pr[B_{ef}|\sigma] \le \frac{p}{2}2^{-n+1}(1-p)^j 2^{-n+1+i}\left(\frac{1+p}{2}\right)^i$$

Lets take the factors one at a time. Firstly, v itself must start blue and turn red. Secondly, all other $v' \in f$ must start blue. Thirdly, all $v' \in f$ coming before v must have second coin tails. Fourthly, all $v' \in e$ coming after v must start red (since v is the last point of e to change color). Finally, all $v' \in e$ coming before v must either start red or start blue and turn red. [The final factor may well be a substantial overestimate. Those $v' \in e$ coming before v which start blue must not only have second coin heads but must themselves lie in an $e' \in H$ monochromatic under the first coloring. Attempts to further improve bounds on $m(n)$ have often centered on this overestimate but (thus far!) to no avail.]

We can then write

$$\Pr[B_{ef}] \leq 2^{1-2n} pE\left[(1+p)^i(1-p)^j\right]$$

where the expectation is over the uniform choice of σ. The following gem therefore completes the argument.

Lemma 5.3: $E\left[(1+p)^i(1-p)^j\right] \leq 1$.

Proof: Fix a matching between $e - \{v\}$ and $f - \{v\}$, think of Mr. & Mrs. Jones; Mr. & Mrs. Smith, etc. Condition on how many of each pair (two Joneses, one Smith, no Taylors,...) come before v. The conditional expectation of $(1+p)^i(1-p)^j$ splits into factors for each pair. When there is no Taylor there is no factor. When there are two Joneses there is a factor $(1+p)(1-p) \leq 1$. When there is one Smith the factor is equally likely to be $1+p$ or $1-p$ and so the conditional expectation gets a factor of one. All factors are at most one so their product is at most one. □

The proof, as it must consider all ways e can be red in the final coloring, may obscure the main obstacle. Fix e, f overlapping in a single vertex v. Let B^*_{ef} be the event that in the first coloring f was all blue and e, except for v, was all red and that v was recolored. There is probability 2^{1-2n} that the first coloring is just so and probability at roughly $1/n$ that v is recolored – since almost surely some $w \in f$ has second coin tails and v is equally likely to be the first such. Then $\sum \Pr[B^*_{ef}]$ is bounded by k^2/n. The subtle analysis of Theorem 5.1 bounds the full $\Pr[B_{ef}]$ by $k^2 p$ - an additional factor of only $\log n$. An improvement in the lower bound of $m(n)$ beyond $2^n n^{1/2}$ would have to, so it seems, find an algorithm that makes the events B^*_{ef} less likely.