Randomized Algorithms

Probabilistic method II: conditional expectations

Alex Gittens
Department of Computer Science, Rensselaer Polytechnic Institute

Fall 2026

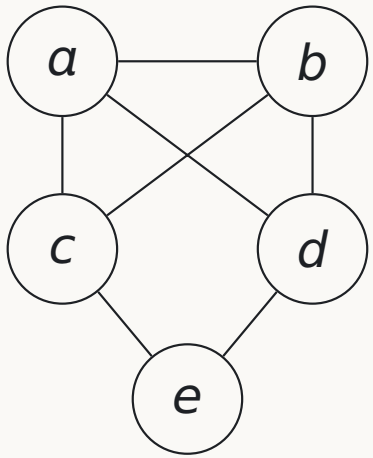
PART 1

MAX-CUT: finding a good cut

1

Finding a large cut

For a graph $G = (V, E)$ with m edges, put each vertex independently into S or T, with equal probability.



Let X be the number of edges crossing between S and T.

$$X = \sum_{uv \in E} \mathbf{1}\{u, v \text{ on opposite sides}\}.$$

Every edge crosses with probability $1/2$, so

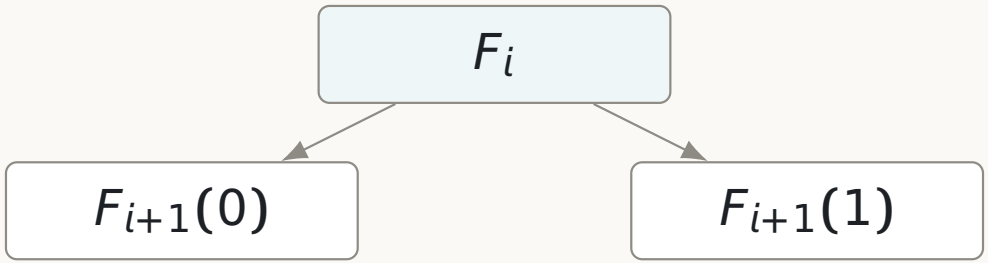
$$\boxed{\mathbb{E}X = m/2.}$$

Some cut therefore has at least $m/2$ crossing edges. How can we find one?

Conditional expectation after another choice

Let $R_1, \ldots, R_n$ be independent fair choices. Once $R_1 = r_1, \ldots, R_i = r_i$ are fixed, write $F_i = \mathbb{E}[X \mid R_{1:i} = r_{1:i}]$.
The *law of total expectation* (the conditional tower property) says

$$\boxed{F_i = \tfrac{1}{2}F_{i+1}(r_{1:i}, 0) + \tfrac{1}{2}F_{i+1}(r_{1:i}, 1).}$$



One of the two children has conditional expectation at least F_i. This remains true after any sequence of earlier choices.

What happens if we fix the first vertex?

Order the vertices $v_1, \ldots, v_n$. We want to choose their sides without using randomness.
Let X still denote the number of crossing edges when all unassigned vertices are placed randomly.

$$\mathbb{E}X = \tfrac{1}{2}\mathbb{E}[X \mid v_1 \in S] + \tfrac{1}{2}\mathbb{E}[X \mid v_1 \in T].$$

At least one of these two conditional expectations is at least $m/2$.
Fix v_1 on whichever side has the larger conditional expectation. The expected number of crossing edges has not decreased.

Can we make the second choice the same way?

Repeating the argument

After fixing $v_1, \ldots, v_i$, let

$$F_i = \mathbb{E}[X \mid \text{the choices for } v_1, \ldots, v_i].$$

The next vertex is equally likely to go to either side, so

$$F_i = \frac{1}{2} F_{i+1}(v_{i+1} \in S) + \frac{1}{2} F_{i+1}(v_{i+1} \in T).$$

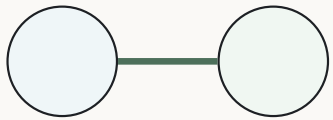
Choose a side for which $F_{i+1} \geq F_i$. Repeating this yields

$$\underbrace{F_0}_{m/2} \leq F_1 \leq \cdots \leq F_n.$$

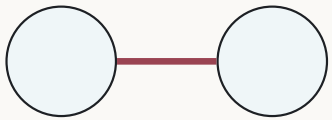
Once every vertex is assigned, there is no randomness left: $F_n = X$.
Hence the final cut has at least $m/2$ crossing edges.
The remaining question is whether we can compute and compare the two conditional expectations efficiently.

How much does an edge contribute?

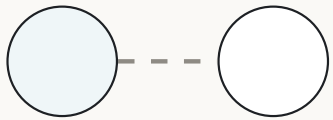
Suppose some vertices have already been assigned. An edge has three possibilities.



opposite sides
1



same side
0



unassigned endpoint
1/2

Let D be the number of already-decided crossing edges, and U the number of edges with at least one unassigned endpoint. By linearity of expectation,

$$\boxed{F_i = D + \frac{1}{2}U.}$$

We can evaluate this quantity from a partial assignment, without enumerating any complete cuts.

Which side should receive the next vertex?

Let v be the next vertex. Write

$$d_S(v) = \#\{\text{assigned neighbors of } v \text{ in } S\},$$
$$d_T(v) = \#\{\text{assigned neighbors of } v \text{ in } T\}.$$

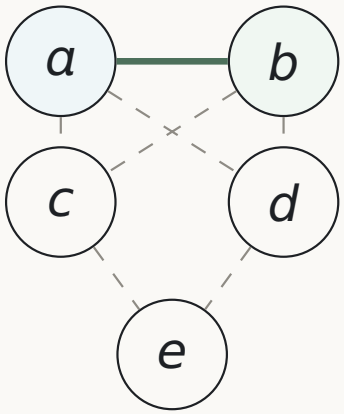
If we put v into S, then $d_T(v)$ edges to assigned vertices cross. If we put it into T, then $d_S(v)$ cross.
Edges from v to unassigned vertices still contribute $1/2$ either way. Every other edge is unchanged. Therefore

$$\boxed{F_{i+1}(v \in S) - F_{i+1}(v \in T) = d_T(v) - d_S(v).}$$

Put v opposite the larger number of its already-assigned neighbors. Break ties arbitrarily.

Choosing where to put a and b

Use the graph with edges $ab, ac, ad, bc, bd, ce, de$, and assign vertices in the order a, b, c, d, e.



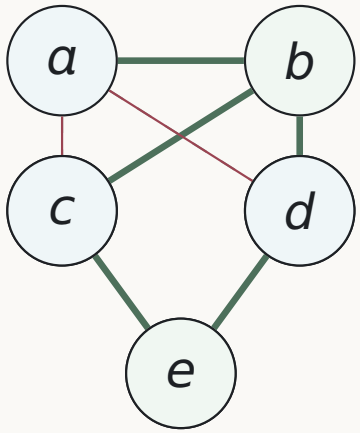
Initially, $F_0 = 7/2$.
Assign a to S. No edge is decided yet, so $F_1 = 7/2$.
For b, the only assigned neighbor is $a \in S$.

$$F_2(b \in S) = 3, \qquad F_2(b \in T) = 4.$$

Choose $b \in T$. Now $F_2 = 1 + 6(1/2) = 4$. The solid green edge is determined to cross; the six dashed edges remain undecided.

Assigning the remaining vertices

We have $a \in S$ and $b \in T$, with $F_2 = 4$.



Put $c \in S$: one previously assigned neighbor on either side. $F_3 = 4$.
Put $d \in S$: again a tie. $F_4 = 4$.
Put $e \in T$: its two assigned neighbors are in S, and $F_5 = 5$.

The final cut has five crossing edges, exceeding the guaranteed $7/2$.

Deterministic MAX-CUT

Input: An undirected graph $G = (V, E)$ with vertices ordered $v_1, \ldots, v_n$.
Output: A partition (S, T) with at least $|E|/2$ crossing edges.

1. Initialize $S = T = \emptyset$.
2. For $i = 1, \ldots, n$, count the already-assigned neighbors of v_i in S and in T.
3. If $d_T(v_i) \geq d_S(v_i)$, put v_i in S. Otherwise put it in T.
4. Return (S, T).

With adjacency lists, checking the neighbors of v_i takes $O(1 + \deg(v_i))$ time. Summing over all vertices gives
$O(n + \sum_i \deg(v_i)) = O(n + 2m) = O(n + m)$.
The guarantee comes from maintaining $F_i \geq F_0 = m/2$.

The method of conditional expectations

More generally, let $R_1, \ldots, R_n$ be independent fair bits and let $X = f(R_1, \ldots, R_n)$. Suppose

$$\mathbb{E}X \geq t.$$

After fixing a prefix $r_1, \ldots, r_i$, set

$$F_i(r_1, \ldots, r_i) = \mathbb{E}[X \mid R_1 = r_1, \ldots, R_i = r_i].$$

Because

$$F_i = \tfrac{1}{2}F_{i+1}(r_1, \ldots, r_i, 0) + \tfrac{1}{2}F_{i+1}(r_1, \ldots, r_i, 1),$$

we can always fix the next bit without decreasing F. At the end,

$$f(r_1, \ldots, r_n) = F_n \geq F_0 \geq t.$$

Computational requirement: We must be able to compare the two child conditional expectations efficiently.

The same method handles minimization

Suppose a randomized construction has a nonnegative cost Y with

$$\mathbb{E}Y \le t.$$

At each step choose the next bit giving the *smaller* conditional expectation. At least one child is no larger than the parent.

$$Y(r_1, \ldots, r_n) \le \mathbb{E}Y \le t.$$

This version will be useful when Y counts violated constraints.

PART 2

MAX-3SAT: satisfying many clauses

2

How many clauses can we satisfy?

A clause is an OR of literals, such as $(x_1 \vee \neg x_2 \vee x_3)$. A 3-CNF formula is an AND of clauses.

Clause	Under $(x_1, x_2, x_3) = (0, 1, 1)$
$(x_1 \vee \neg x_2 \vee x_3)$	satisfied
$(\neg x_1 \vee x_2 \vee \neg x_3)$	satisfied
$(x_1 \vee \neg x_2 \vee \neg x_3)$	falsified

We may not be able to satisfy every clause. *MAX-3SAT* asks for an assignment satisfying as many as possible.
For a clause on three distinct variables, a random assignment falsifies all three literals with probability $1/8$. Thus every clause is satisfied with probability $7/8$. If X counts satisfied clauses among m clauses,

$$\boxed{\mathbb{E}X = 7m/8.}$$

Can we achieve this guarantee with deterministic choices?

Probability that a clause is satisfied

After fixing some variables, define for each clause C

$$p_C = \Pr(C \text{ is satisfied} \mid \text{current assignments}).$$

$$(x_1 \ \vee \ \neg x_2 \ \vee \ x_3)$$

Fixed variables	Remaining question	p_C
none	Are any of three literals true?	7/8
$x_1 = 0$	Will $\neg x_2$ or x_3 be true?	3/4
$x_1 = 0,\ x_2 = 1$	Will x_3 be true?	1/2
$x_1 = 0,\ x_2 = 1,\ x_3 = 0$	All three are false	0

If any fixed literal is true, $p_C = 1$. Otherwise, with r distinct unassigned variables, $p_C = 1 - 2^{-r}$. If X counts satisfied clauses, then $\mathbb{E}[X \mid \text{current assignments}] = \sum_C p_C$.

Deterministic MAX-3SAT

Input: A 3-CNF formula with n variables and m clauses, each using three distinct variables.
Output: An assignment satisfying at least $7m/8$ clauses.

1. Set $i = 1$ and begin with no assigned variables.
2. For $b = 0, 1$, temporarily set $x_i = b$ and sum the resulting conditional satisfaction probabilities p_C over all clauses. Call the sum $F(b)$.
3. Permanently set x_i to whichever b gives the larger $F(b)$ (ties arbitrary).
4. Increment i and repeat Steps 2–3 until all variables are assigned. Return the assignment.

Guarantee: Each step preserves the conditional expected number of satisfied clauses, initially $7m/8$. At the end it equals the actual number satisfied. **Runtime:** Each of n rounds evaluates two clause sums, each in $O(m)$ time because every clause has three literals. Total $O(nm)$.

PART 3

Set balancing: avoiding bad events

3

Balancing overlapping sets

Suppose we want to divide people into two teams, while giving several overlapping groups roughly equal representation on each team.
Let $S_1, \ldots, S_M \subseteq [n]$, each of size at most k. Assign each item a sign $\chi_i \in \{-1, +1\}$. For each set, define the imbalance

$$z_j = \sum_{i \in S_j} \chi_i.$$

	1	2	3	4	5	6	7	z_j
S_1	•	•		•			•	2
S_2		•	•	•	•			0
S_3	•		•			•	•	2
χ_i	+	−	+	+	−	−	+	

Each column affects every set containing that item. Can we keep all $|z_j|$ small simultaneously?

Random signs give a simultaneous bound

Choose all signs independently and uniformly. Hoeffding's inequality gives, for each j,

$$\Pr(|Z_j| \geq t) \leq 2\exp\left(-\frac{t^2}{2k}\right).$$

By the union bound,

$$\Pr(\exists j : \ |Z_j| \geq t) \leq 2M\exp\left(-\frac{t^2}{2k}\right).$$

Taking

$$\boxed{t = \sqrt{2k\ln(4M)}}$$

makes the right side at most $1/2$. Thus some signing satisfies $|Z_j| < t$ for every j.

How do we find such signs without trying exponentially many assignments?

Count the violated constraints

Let $A_j = \{|Z_j| \geq t\}$ and let

$$Y = \sum_{j=1}^{M} \mathbf{1}_{A_j}$$

be the number of sets whose imbalance is too large.
By linearity of expectation and the same tail bounds,

$$\mathbb{E}Y = \sum_{j=1}^{M} \Pr(A_j) \leq \frac{1}{2} < 1.$$

Since Y is a nonnegative integer, some complete signing must have $Y = 0$. Can we find such a signing?

Finding a signing with no violations

After fixing a prefix $s_{1:i}$, define

$$H_i(s_{1:i}) = \mathbb{E}[\,Y \mid \chi_{1:i} = s_{1:i}\,].$$

The next sign is fair, so the tower property gives

$$H_i = \tfrac{1}{2} H_{i+1}(s_{1:i}, +1) + \tfrac{1}{2} H_{i+1}(s_{1:i}, -1).$$

One choice keeps H from increasing. Repeating gives

$$Y = H_n \le H_0 = \mathbb{E}Y \le \frac{1}{2}.$$

Since Y is an integer, the final signing violates no constraints. **But how do we evaluate the two choices of H_{i+1}?**

What must we compute?

Recall $Y = \sum_{j=1}^{M} \mathbf{1}_{A_j}$. By linearity of conditional expectation,

$$\boxed{H_i(s_{1:i}) = \sum_{j=1}^{M} \Pr(A_j \mid s_{1:i}).}$$

To evaluate a proposed next sign, we must calculate the probability that *each individual set* exceeds its imbalance threshold, conditioned on the signs already assigned.
The sets can overlap and their violation events can be dependent. This does not affect the sum, because expectation is linear.
How do we compute $\Pr(A_j \mid s_{1:i})$ for one set?

What remains random for one set?

After fixing $\chi_1, \ldots, \chi_i$, separate the terms of one set's imbalance:

$$Z_j = a_j + \sum_{\ell=1}^{r_j} \xi_\ell,$$

where a_j is the sum of the signs already fixed in S_j, and r_j is the number of its still-unassigned items.

Item	1	2	3	4	5	6
In S_j?	•	•		•	•	•
Sign	+	−	+	?	?	?

Here $a_j = (+1) + (-1) = 0$ and $r_j = 3$. The random signs outside S_j do not affect its imbalance.

From random signs to a binomial variable

For a set with r_j still-random signs, let K be the number of them equal to $+1$.
Then $K \sim \mathrm{Binomial}(r_j, 1/2)$, since the remaining signs are independent and fair. There are K positive signs and $r_j - K$ negative signs, so their sum is

$$K - (r_j - K) = 2K - r_j.$$

Thus

$$\boxed{Z_j = a_j + 2K - r_j.}$$

For example, if $r_j = 3$ and $K = 2$, the remaining signs sum to $+1$. The entire set has imbalance $a_j + 1$.

Probability that one set violates the bound

The event A_j is $\{|Z_j| \geq t\}$. For a fixed prefix, the binomial representation gives

$$\Pr(A_j \mid s_{1:i}) = \Pr(|a_j + 2K - r_j| \geq t).$$

The possible values of K are $0, \ldots, r_j$, with probabilities $2^{-r_j}\binom{r_j}{q}$. Hence

$$\Pr(A_j \mid \chi_{1:i}) = 2^{-r_j} \sum_{q=0}^{r_j} \binom{r_j}{q} \mathbf{1}\{|a_j + 2q - r_j| \geq t\}.$$

We can evaluate this probability with at most $k+1$ terms because $r_j \leq k$.

Deterministic set balancing

Input: Sets $S_1, \ldots, S_M \subseteq [n]$ of size at most $k \geq 1$.
Output: Signs $s_1, \ldots, s_n \in \{-1, +1\}$ with every imbalance below $t = \sqrt{2k \ln(4M)}$.

1. Compute t. Initially no signs are fixed. Set $i = 1$.
2. For $b \in \{-1, +1\}$, temporarily assign $s_i = b$. For every set, compute its conditional violation probability using the binomial formula, then sum these M probabilities to obtain $H(b)$.
3. Fix s_i to whichever b gives the smaller $H(b)$.
4. Increment i and repeat Steps 2–3 until all signs are fixed. Return the signs.

Guarantee: The returned signing violates no constraints:

$$\left| \sum_{i \in S_j} s_i \right| < \sqrt{2k \ln(4M)} \qquad \text{for all } j.$$

Cost of computing the next sign

At one step, the algorithm evaluates two possible sign choices.

- For each choice, it computes M conditional violation probabilities.
- Each probability is a sum of at most $k + 1$ binomial terms.

Each choice takes $O(Mk)$ time, and we make n choices. Thus the running time is

$$\boxed{O(nMk).}$$

When conditional probabilities are expensive

In unweighted set balancing, the remaining signs all have coefficient 1.
The sum depends only on how many are positive, so binomial tails suffice.
Now suppose the constraints are weighted:

$$Z_j = \sum_i a_{ji} \chi_i,$$

where the a_{ji} are integers, possibly of very different sizes.
After fixing a prefix of signs, the remaining random quantity for one constraint is

$$Z_j = b_j + \sum_{\ell \in U_j} a_{j\ell} \xi_\ell.$$

To use conditional expectations directly, we would need

$$\Pr\left(\left| b_j + \sum_{\ell \in U_j} a_{j\ell} \xi_\ell \right| \geq t_j \right).$$

Why is the weighted probability harder?

Take four unassigned signs with coefficients $2, 5, 11, 23$. If precisely two signs are positive, consider these assignments:

ξ_1	ξ_2	ξ_3	ξ_4	$2\xi_1 + 5\xi_2 + 11\xi_3 + 23\xi_4$
$+$	$+$	$-$	$-$	-27
$+$	$-$	$+$	$-$	-15
$-$	$-$	$+$	$+$	27

The number of positive signs no longer determines the weighted sum. We need to know *which* signs are positive. With r unassigned signs, direct enumeration examines 2^r assignments. For binary-encoded integer weights, dynamic programming can depend on their magnitudes rather than just the input length. We therefore want a bound that we can update without computing the exact weighted tail probability.

Pessimistic estimators: the idea

Let Y count violated constraints. Rather than compute $\mathbb{E}[Y \mid s_{1:i}]$ exactly, maintain a computable potential Φ_i on partial sign assignments.
We want three properties:

1. Initially, $\Phi_0 < 1$.
2. For every partial assignment, at least one next-sign choice satisfies $\Phi_{i+1} \leq \Phi_i$.
3. Every bad complete assignment has $\Phi_n \geq 1$.

Choose the next sign that gives the smaller potential. At the end,

$$\Phi_n \leq \Phi_0 < 1,$$

so no constraint is violated.

Exponential potentials from Chernoff bounds

We call such a potential a *pessimistic estimator*: it may overestimate the chance or amount of failure, but its bound is strong enough to certify a good final assignment.
Chernoff's method gives a natural construction. A threshold violation can be upper-bounded by an exponential, for example

$$\mathbf{1}\{Z_j \geq t_j\} \leq e^{-\lambda_j t_j} e^{\lambda_j Z_j}.$$

The key computational advantage is

$$e^{\lambda_j \sum_i a_{ji} x_i} = \prod_i e^{\lambda_j a_{ji} x_i}.$$

For independent remaining signs, the conditional expectation of this product factors.
The appendix derives this estimator for weighted set balancing and proves its guarantee.

What have we used?

In each example, we maintained a quantity defined on partially fixed random choices.

- For MAX-CUT and MAX-3SAT, it was the expected objective value. We kept it from decreasing.
- For set balancing, it was the expected number of violations. We kept it from increasing.

The argument guarantees a good root-to-leaf path. An efficient derandomization also requires that we can evaluate the two choices at each step.
Next: the Lovász Local Lemma, which can establish existence even when the global union bound is too weak.

PART 4

Appendix: exponential pessimistic estimators

Weighted balancing

Suppose instead that each constraint is a weighted signed sum:

$$Z_j = \sum_{i=1}^{n} a_{ji} x_i, \qquad x_i \in \{-1, +1\}.$$

Write $L_j^2 = \sum_i a_{ji}^2$ and assume $L_j > 0$.
Random signs satisfy the Hoeffding bound

$$\Pr(|Z_j| \geq t_j) \leq 2e^{-t_j^2/(2L_j^2)}.$$

Choosing $t_j = L_j\sqrt{2\ln(4M)}$ makes the sum of these bounds at most $1/2$. A simultaneous good signing exists.
But after fixing some signs, evaluating the exact probability that a weighted sum crosses its threshold may require tracking exponentially many distinct subset sums. Can we avoid those exact probabilities?

An upper bound that factors

For any $\lambda_j > 0$,

$$\mathbf{1}\{|Z_j| \geq t_j\} \leq e^{-\lambda_j t_j}\left(e^{\lambda_j Z_j} + e^{-\lambda_j Z_j}\right) =: W_j.$$

Let

$$W = \sum_{j=1}^{M} W_j, \qquad Y = \sum_{j=1}^{M} \mathbf{1}\{|Z_j| \geq t_j\}.$$

Then $Y \leq W$ for every complete signing.
The advantage of exponentials is that

$$e^{\lambda_j Z_j} = \prod_i e^{\lambda_j a_{ji} \chi_i},$$

so expectations over independent, unassigned signs factor into products.

The conditional exponential potential

After fixing signs $s_1, \ldots, s_i$, write

$$b_j = \sum_{\ell \le i} a_{j\ell} s_\ell .$$

For an unassigned fair sign χ_ℓ,

$$\mathbb{E}[\, e^{\lambda_j a_{j\ell} \chi_\ell} \,] = \cosh(\lambda_j a_{j\ell}).$$

Therefore the conditional expectation of W is

$$\boxed{\Phi_i = \sum_{j=1}^{M} 2 e^{-\lambda_j t_j} \cosh(\lambda_j b_j) \prod_{\ell > i} \cosh(\lambda_j a_{j\ell}).}$$

Both candidate values of Φ_{i+1} can be evaluated using the coefficients, without calculating any weighted-sum tail probabilities.

The potential starts below one

Use $\lambda_j = t_j / L_j^2$. Since $\cosh u \le e^{u^2/2}$,

$$\begin{aligned}
\Phi_0 &= \sum_{j=1}^{M} 2e^{-\lambda_j t_j} \prod_i \cosh(\lambda_j a_{ji}) \\
&\le \sum_{j=1}^{M} 2 \exp\left(-\lambda_j t_j + \frac{\lambda_j^2 L_j^2}{2} \right) \\
&= \sum_{j=1}^{M} 2e^{-t_j^2/(2L_j^2)} \\
&= \frac{1}{2} < 1.
\end{aligned}$$

The last equality uses $t_j = L_j \sqrt{2 \ln(4M)}$.

Choosing signs with the surrogate potential

Because $\Phi_i = \mathbb{E}[\,W \mid s_{1:i}]$,

$$\Phi_i = \tfrac{1}{2}\Phi_{i+1}(s_{1:i}, +1) + \tfrac{1}{2}\Phi_{i+1}(s_{1:i}, -1).$$

Choose the sign giving the smaller next potential. This ensures

$$\Phi_n \le \Phi_0 \le \frac{1}{2}.$$

At a complete assignment, $\Phi_n = W$. If any constraint were violated, its corresponding W_j would be at least 1, giving $W \ge 1$, a contradiction. Thus all weighted constraints satisfy

$$\boxed{|Z_j| < L_j\sqrt{2\ln(4M)} \qquad (1 \le j \le M).}$$

Pessimistic estimators

The weighted construction illustrates the general idea. Instead of evaluating the exact probability of failure, maintain a potential on partial assignments with three properties:

1. It starts below 1.
2. One choice of the next random variable does not increase it.
3. Every bad complete assignment has potential at least 1.

Here the potential was the conditional expectation of an exponential upper bound on the number of bad constraints. The bound can be computed even when exact weighted tail probabilities are inconvenient.