Friendly bisections of random graphs

Ashwin Sah

Massachusetts Institute of Technology

Joint work with Asaf Ferber (UCI), Matthew Kwan (IST Austria), Bhargav Narayanan (Rutgers), and Mehtaab Sawhney (MIT)

October 26, 2021

Unfriendly cuts

Definition

A *cut* of a graph is a partition of its vertices into nonempty sets A, B. The *size* of the cut is the number of edges between A and B. A *max (min) cut* is a cut whose size is maximum (minimum) possible.

Definition

The *friendliness* of a vertex v with respect to a cut (A, B) is

$$\Delta_{A,B}(v) = \begin{cases} \deg_A(v) - \deg_B(v) & \text{if } v \in A \\ \deg_B(v) - \deg_A(v) & \text{if } v \in B \end{cases}$$

- In a maximum cut (A, B), every vertex v has $\Delta_{A,B}(v) \leq 0$.
- Thus every graph has an unfriendly cut.

Friendly cuts

- A minimum cut need not be friendly. In fact, there are graphs where every nontrivial cut (A, B) satisfies $\Delta_{A,B}(v) < 0$ for some v.

Theorem (Stiebitz (1996))

In any graph G there is a cut (A, B) for which $\Delta_{A,B}(v) \geq -1$ for all vertices v.

Conjecture (Ban and Linial (2016))

For any d, there are finitely many d-regular graphs with no friendly cut.

Theorem (Linial and Louis (2020))

Fix even d. A random d-regular graph on n vertices has a friendly cut whp.

Bisections

- A max bisection cannot have $a \in A$ and $b \in B$ with $\Delta_{A,B}(a), \Delta_{A,B}(b) > 0$.
- Key issue: one side could be all friendly, and the other all unfriendly.

Conjecture (Füredi)

In $\mathbb{G}(n, 1/2)$, whp there is a bisection (A, B) with $\Delta_{A,B}(v) \geq 0$ for all but $o(n)$ vertices v.

Question (Bollobás and Scott (2002))

In any graph G, is there a bisection (A, B) where $\Delta_{A,B}(v) \leq 1$ for all v?

- In the case of bisections, unfriendly and friendly are quite similar, considering the complement graph (though off-by-one).

Theorem (Ferber, Kwan, Narayanan, S., and Sawhney (2021+))

In $\mathbb{G}(n, 1/2)$ a bisection with $\Delta_{A,B}(v) \geq 0$ for $(1 - o(1))n$ v exists whp.

A simple idea

- Considering a minimum bisection, similar to earlier, shows the following.

Lemma

There is a bisection for which all $a \in A$ satisfy $\Delta_{A,B}(a) \geq -1$ or all $b \in B$ satisfy $\Delta_{A,B}(b) \geq -1$.

- How do we rule out the “conspiracy” that only one side is friendly, and the other is partially (or even fully) unfriendly?
- Idea: start with a random bisection and then repeatedly choose *random* unfriendly vertices $a \in A$ and $b \in B$ to swap.
- This eventually ends in a bisection where all of one side is fully friendly.
- Let X_t be the number of unfriendly vertices in A after t swaps. If we can show $\operatorname{Var} X_t = o(n^2)$ for all $t \leq n^2$, we are done. (Why?)

Try 1: Efron–Stein concentration

Theorem (Efron–Stein)

Let $Z = f(X_1, \ldots, X_N)$ be a function of independent random variables $X_1, \ldots, X_N$, and let $Z^{(i)}$ be the result of resampling X_i. Then $\operatorname{Var} Z \le \sum_{i=1}^{N} \mathbb{E}(Z - Z^{(i)})^2$.

- Unfortunately, it is hard to rule out a "contagion" where resampling a single edge completely changes the result of the process.

Try 2: Differential equations method

(1) Find an ensemble of random variables such that the expected change in each random variable at step t can be expressed/approximated in terms of the variables in the ensemble at step t.

(2) Solve a system of differential equations to obtain an ansatz for the trajectory of the process.

(3) Use martingale concentration to prove the process concentrates around this trajectory.

- This has seen use in analyzing random graph processes such as the triangle-removal process.
- However, it relies heavily on injecting randomness into each new step of the process.
- We have some randomness at each step, but the information revealed to understand the next step is essentially the degrees of all the vertices, which within $O(n)$ steps quickly amounts to revealing most of $\mathbb{G}(n, 1/2)$.

Idea 1: "nibbling" the process

Lemma

In a random graph whp every bisection (A, B) satisfies: (a) all but εn v have $|\Delta_{A,B}(v)| = \Omega_\varepsilon(\sqrt{n})$, and (b) $\sum_{v \in A \cup B} \Delta_{A,B}(v) = O(n^{3/2})$ whp.

Proof sketch of (a).

For fixed A, B, v we see $\mathbb{P}[|\Delta_{A,B}(v)| \le c\sqrt{n}] = O(c)$. Therefore if $c = \exp(-\Omega(1/\varepsilon))$, then a set S of more than εn vertices satisfy this with probability $\le (O(c))^{\Omega(\varepsilon n)} \le 16^{-n}$. Now take a union over A, B, S. □

- Now, if there are many unfriendly vertices on both sides, take a batch of εn vertices on both sides (most of which are *very* unfriendly by (a)).
- Swap them all at once: after $\exp(O(1/\varepsilon))$ swaps, one of the sides must have almost all vertices friendly.
- Thus we only need to reveal information to analyze $O(1)$ steps!

Idea 2: Iterative exposure

- Given bisection (A, B), if we want to know which vertices are most unfriendly, we need to know the degree of each vertex v into A and B.
- We reveal precisely this information at each step and nothing else.
- That is, if the bisection at step t is (A_t, B_t), reveal $\deg_{A_t}(v)$ and $\deg_{B_t}(v)$ for all v, *and nothing else*.
- To analyze the first step we have a random graph with *constrained degrees* to two parts.
- To analyze the second we have a random graph with constrained degrees into four parts.
- $\cdots$
- Such random graph models can be studied in our dense setting using graph enumeration results of McKay and Wormald (1990) and Canfield, Greenhill, and McKay (2008), which are derived using saddle-point analysis of contour integrals of generating functions. (Other regimes form a rich line of research.)

Idea 3: concentration of empirical degree distributions

- For each vertex, record a "history" of which side it has been on at each step, and partition vertices according to their histories.
- Record the entire empirical joint distribution of degrees in each part to the tuple of all parts, e.g., for each j record

$$\mu_i = \mathrm{Unif}\left\{\left(\frac{\deg_{P_j} v - |P_j|/2}{\sqrt{n}}\right)_{j\in I} : v \in P_i\right\}.$$

- In principle, one could show that these distributions are close to certain continuous distributions at all steps (at the very start, we have two parts P_1, P_2 and μ_1, μ_2 are both two-dimensional Gaussians).
- In practice, this is very challenging as the formulas become exceedingly complex past one stage, among other technical difficulties.
- Instead, we simply show this entire empirical distribution concentrates around some "ideal" distribution abstractly using the second-moment method, without actually computing it.

Implementation

We inductively guarantee the following at each step k:

(1) Each of the 2^k possible parts P has size concentrated near $\alpha_P n$

(2) For each part P there is some distribution $\mathcal{L}_k^{(P)}$ such that the 2^k-dimensional empirical degree distribution for part P concentrates around $\mathcal{L}_k$ (in some metric on distributions).

(3) $\mathcal{L}_k$ is anticoncentrated (i.e., the distribution of degrees is "smooth").

(4) $\mathcal{L}_k$ has sub-Gaussian tails (outlier degrees are unlikely).

- To run the next step, we reveal the degrees of all vertices to all the parts satisfying the above guarantees (which we inductively know occurs whp).
- Then, using the remaining randomness in this degree-constrained model (computed using graph enumeration results), show that (a) the number of vertices in each new part concentrates, (b) the new empirical degree distributions concentrate, and (c) the new data is well-behaved.

Further directions: Majority dynamics

Definition

The majority dynamics on G partitioned into R, B is as follows: at each step, swap all vertices with strictly more neighbors on the other side.

Conjecture (Tran and Vu (2019))

Let $\mathbb{G}(2n+1, 1/2)$ be partitioned into R, B of sizes $n+1, n$. Then for some $\varepsilon > 0$, with probability $\geq 1/2 + \varepsilon$, as $n \to \infty$ the majority dynamics ends with $R = [2n+1]$.

Theorem (S. and Sawhney (2021+))

In the above, $R = [2n+1]$ occurs within 3 steps with probability $\mathbb{P}[\mathcal{N}(0,1) \leq \sqrt{2/\pi}] + O(n^{-c})$.

- Idea: we can analyze the steps of this process using iterative degree revelation. The "leader" after one day ends up "winning".

Further directions

- In a random graph, can we find a bisection where *every* vertex is friendly?
- Ban–Linial conjecture: there are finitely many d-regular graphs with no friendly cut.
- Bollobás–Scott conjecture: every graph has a bisection (A, B) with $\Delta_{A,B}(v) \leq 1$ for all v.
- More applications of this technique of iterated degree revelation (friendly bisections, majority dynamics, ?).