Bad Tournaments

Eugene Francisco
advised by Maya Sankar

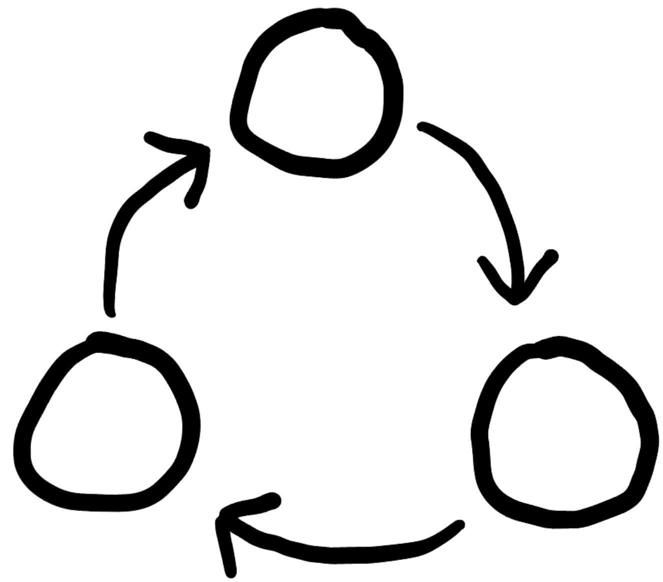

Tournaments

- A $tournament$ on n players T_n is a type of directed graph.
- **Every vertex** has **exactly one** directed edge to **every other** (distinct) vertex. So either $v \to u$ or $u \to v$ but not both.

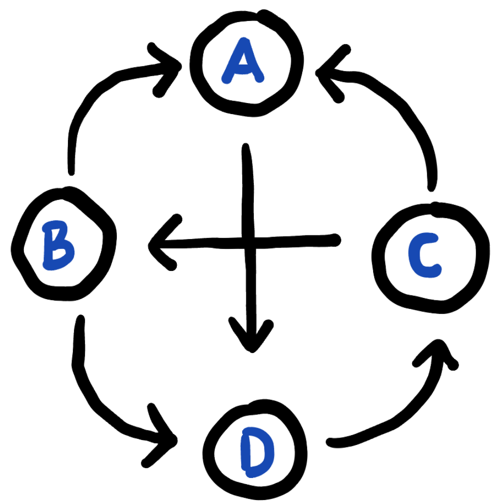

Tournaments

- A *tournament* on n players T_n is a type of directed graph.
- **Every vertex** has **exactly one** directed edge to **every other** (distinct) vertex. So either $v \to u$ or $u \to v$ but not both.
- If $v \to u$, then v "beats" u.

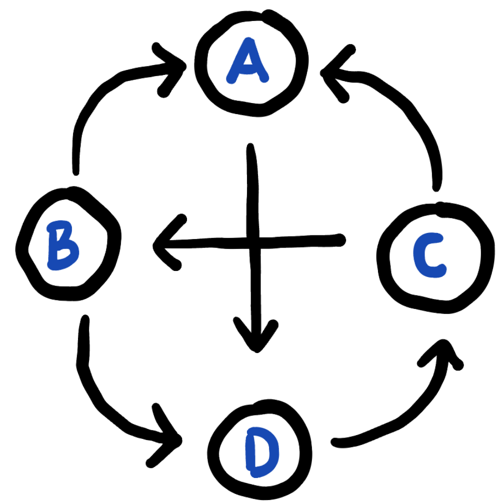

Tournaments

- A *tournament* on n players T_n is a type of directed graph.
- **Every vertex** has **exactly one** directed edge to **every other** (distinct) vertex. So either $v \to u$ or $u \to v$ but not both.
- If $v \to u$, then v "beats" u.

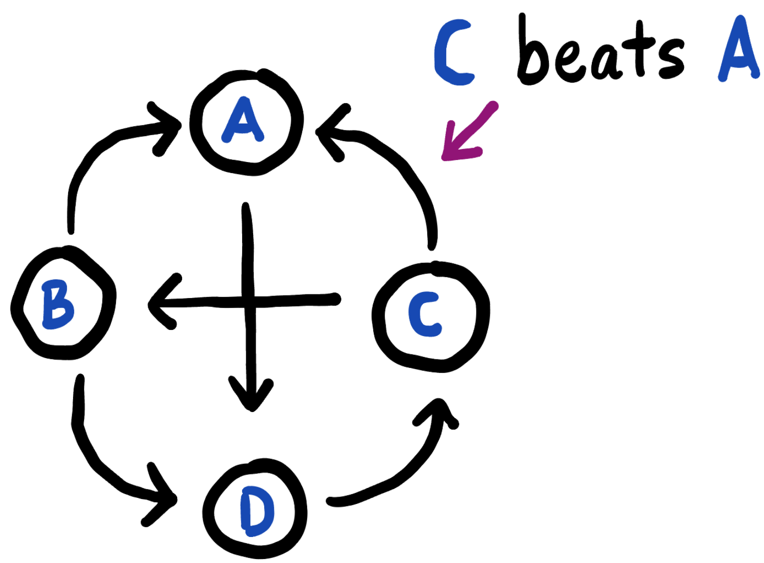

Tournaments

- A *tournament* on n players T_n is a type of directed graph.
- **Every vertex** has **exactly one** directed edge to **every other** (distinct) vertex. So either $v \to u$ or $u \to v$ but not both.
- If $v \to u$, then v "beats" u.

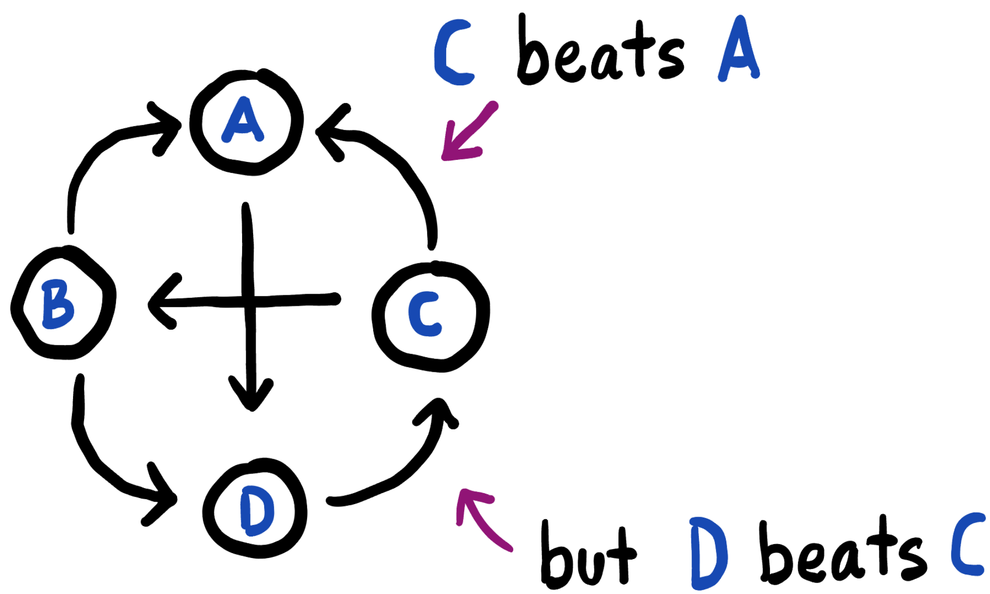

Rankings

- A ranking π is just an ordering of the vertices.

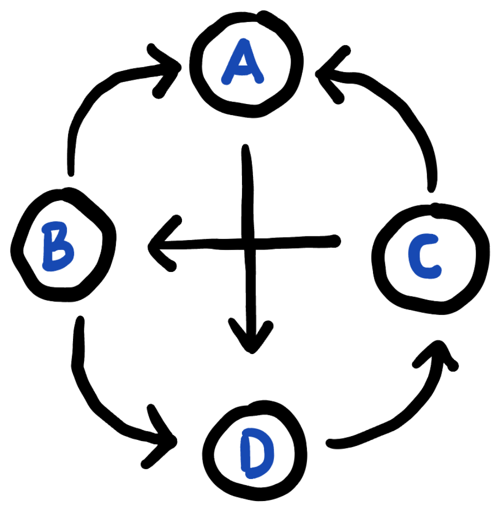

Rankings

- A ranking π is just an ordering of the vertices.
- $(A > B > C > D)$ is one possible ranking.
- So is $(B > D > A > C)$.

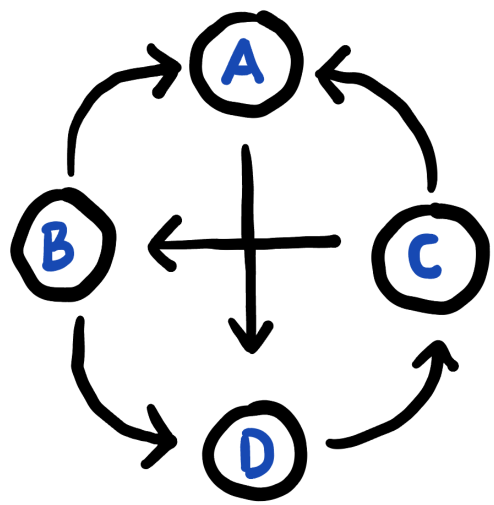

Consistencies

- Let T_n be a tournament and π a ranking of the tournament.

- If $v \to u$ in the tournament, and v before u in π, then the edge (v, u) is *consistent* with π.

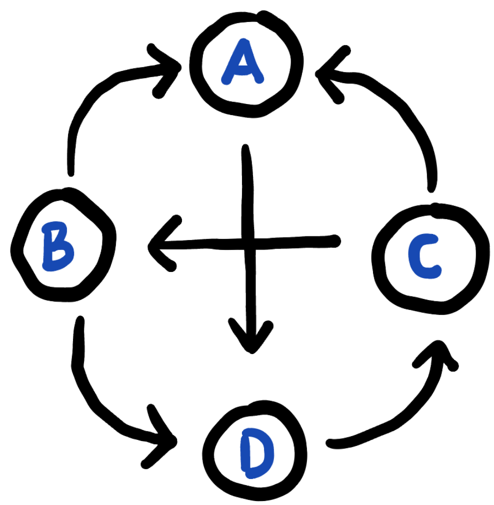

Consistencies

- Let T_n be a tournament and π a ranking of the tournament.

- If $v \to u$ in the tournament, and v before u in π, then the edge (v, u) is *consistent* with π.

- If $\pi = (A > B > C > D)$, then the pair (B, D) is consistent, because $B \to D$ in the graph.

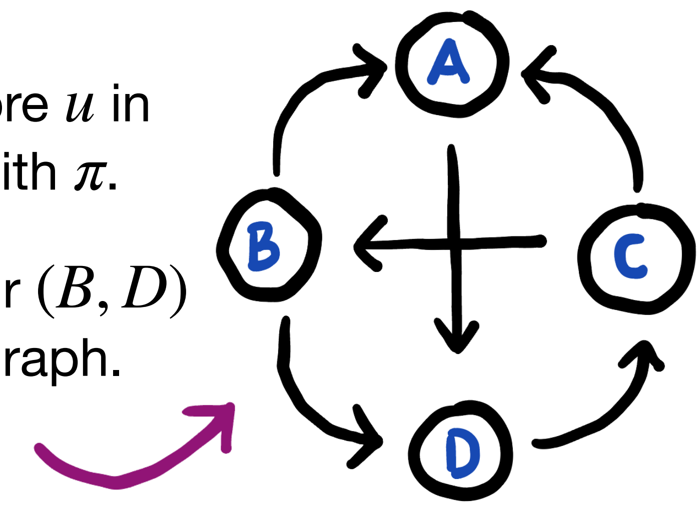

A Best Ranking?

Let $c(\pi, T_n)$ be the number of edges in T_n that are consistent with π.

And let $c(T_n) = \max_{\pi} c(\pi, T_n)$ be the most consistencies any of the rankings have.

A Best Ranking?

Let $c(\pi, T_n)$ be the number of edges in T_n that are consistent with π.

And let $c(T_n) = \max_{\pi} c(\pi, T_n)$ be the most consistencies any of the rankings have.

$$\boxed{\text{We claim } c(T_n) \geq \frac{1}{2}\binom{n}{2}}$$

In English: "For any tournament you give me, I can find a a ranking that gives better than $\frac{1}{2}\binom{n}{2}$ consistencies."

$$c(T_n) = \max_{\pi} c(\pi, T_n)$$

A Best Ranking?

We claim $\boxed{c(T_n) \geq \frac{1}{2}\binom{n}{2}}$

Proof: If $\pi = (\pi_1, \ldots, \pi_n)$, and (π_i, π_j) is inconsistent in π, then (π_i, π_j) must be consistent in $\pi' = (\pi_n, \pi_{n-1}, \ldots, \pi_1)$.

$$c(T_n) = \max_{\pi} c(\pi, T_n)$$

A Best Ranking?

$\boxed{\text{We claim } c(T_n) \geq \frac{1}{2}\binom{n}{2}}$

Proof: If $\pi = (\pi_1, \ldots, \pi_n)$, and (π_i, π_j) is inconsistent in π, then (π_i, π_j) must be consistent in $\pi' = (\pi_n, \pi_{n-1}, \ldots, \pi_1)$.

So $c(\pi, T_n) + c(\pi', T_n) = \binom{n}{2}$.

$$c(T_n) = \max_{\pi} c(\pi, T_n)$$

A Best Ranking?

We claim $c(T_n) \geq \frac{1}{2}\binom{n}{2}$

Proof: If $\pi = (\pi_1, \ldots, \pi_n)$, and (π_i, π_j) is inconsistent in π, then (π_i, π_j) must be consistent in $\pi' = (\pi_n, \pi_{n-1}, \ldots, \pi_1)$.

So $c(\pi, T_n) + c(\pi', T_n) = \binom{n}{2}$.

So $c(T_n) \geq \frac{1}{2}\binom{n}{2}$.

$\square$

Bad Tournaments

We know $c(T_n) \geq \frac{1}{2}\binom{n}{2}$ for any tournament.

Question: Can we make a tournament where rankings can't get much better than this?

Bad Tournaments

We know $c(T_n) \geq \frac{1}{2}\binom{n}{2}$ for any tournament.

Question: Can we make a tournament where rankings can't get much better than this?

$$\boxed{\textbf{Goal: } c(T_n) \leq \frac{1}{2}\binom{n}{2} + \textit{something small}}$$

In English: "I want a tournament where the best ranking I can manage is only slightly better than what I'm guaranteed."

$$\boxed{\textbf{Goal: } c(T_n) \leq \frac{1}{2}\binom{n}{2} + \textit{something small}}$$

One Construction: Pick a prime p (for technical reasons, $p \equiv 3 \mod 4$). Let the vertices of T_p be $\{1, \ldots, p-1\}$ and have $x \to y$ in T_p precisely when $x - y$ is a square mod p.

Goal: $c(T_n) \le \frac{1}{2}\binom{n}{2} + \textit{something small}$

One Construction: Pick a prime p (for technical reasons, $p \equiv 3 \mod 4$). Let the vertices of T_p be $\{0,\ldots,p-1\}$ and have $x \to y$ in T_p precisely when $x - y$ is a square mod p.

The Big Theorem: $c(T_p) \le \frac{1}{2}\binom{p}{2} + O(p^{3/2}\log p)$

In English: "The *best* ranking I can manage on these tournaments is only slightly better than what I can guarantee."

Adjacency Matrix: Let D be the $p \times p$ matrix defined by

$$D_{ij} = \begin{cases} 1 & i - j \text{ a square,} \\ -1 & i - j \text{ not a square,} \\ 0 & i = j\,. \end{cases}$$

The entries of D tell us what the tournament looks like. D is called the adjacency matrix of our tournament T_p.

First, some setup: Let $A, B \subset \{0, \ldots, p-1\}$ be two subsets of vertices. Define $e(A, B)$ as the number of edges in T_p which go from A to B. By construction,

$$e(A, B) - e(B, A) = \sum_{i \in A} \sum_{j \in B} D_{ij}$$

First, some setup: Let $A, B \subset \{0, \ldots, p-1\}$ be two subsets of vertices. Define $e(A, B)$ as the number of edges in T_p which go from A to B. By construction,

$$e(A,B) - e(B,A) = \sum_{i \in A} \sum_{j \in B} D_{ij}$$

1 when $i \to j$

-1 when $i \leftarrow j$

First, some setup: Let $A, B \subset \{0,\ldots,p-1\}$ be two subsets of vertices. Define $e(A,B)$ as the number of edges in T_p which go from A to B. By construction,

$$e(A,B) - e(B,A) = \sum_{i \in A} \sum_{j \in B} D_{ij}$$

1 when $i \to j$

-1 when $i \leftarrow j$

Lemma 2: $|e(A,B) - e(B,A)| = \left| \sum_{i \in A} \sum_{j \in B} D_{ij} \right| \le \sqrt{|A||B|p}$

$\boxed{\textbf{Lemma 2: } |e(A,B) - e(B,A)| \leq \sqrt{|A||B|p}}$

In English: (Roughly) the difference in the number of edges that go from $A \to B$ against $B \to A$ is at most proportional to the product of their sizes.

$\boxed{\textbf{Lemma 2: } |e(A,B) - e(B,A)| \le \sqrt{|A||B|p}}$

In English: (Roughly) the difference in the number of edges that go from $A \to B$ against $B \to A$ is at most proportional to the product of their sizes.

Proof: We don't have time.

The important thing is that **we get to choose** A **and** B however we like, and we know the number of edges between A and B is **controlled by how many vertices** A **and** B have.

The Big Theorem: $c(T_p) \leq \frac{1}{2}\binom{p}{2} + O(p^{3/2}\log p)$

Proof: Let π be a ranking. Let A_1 be the first half of π and A_2 be the second half of π.

Important: $e(A_1, A_2)$ counts the edges between A_1, A_2 that are consistent with π (!) And $e(A_2, A_1)$ counts the edges between A_1, A_2 consistent with π' !

$$\pi = (\underbrace{\pi_1, \pi_2, \dots, \pi_{\sim p/2}}_{A_1}, \underbrace{\pi_{\sim p/2}, \dots, \pi_p}_{A_2})$$

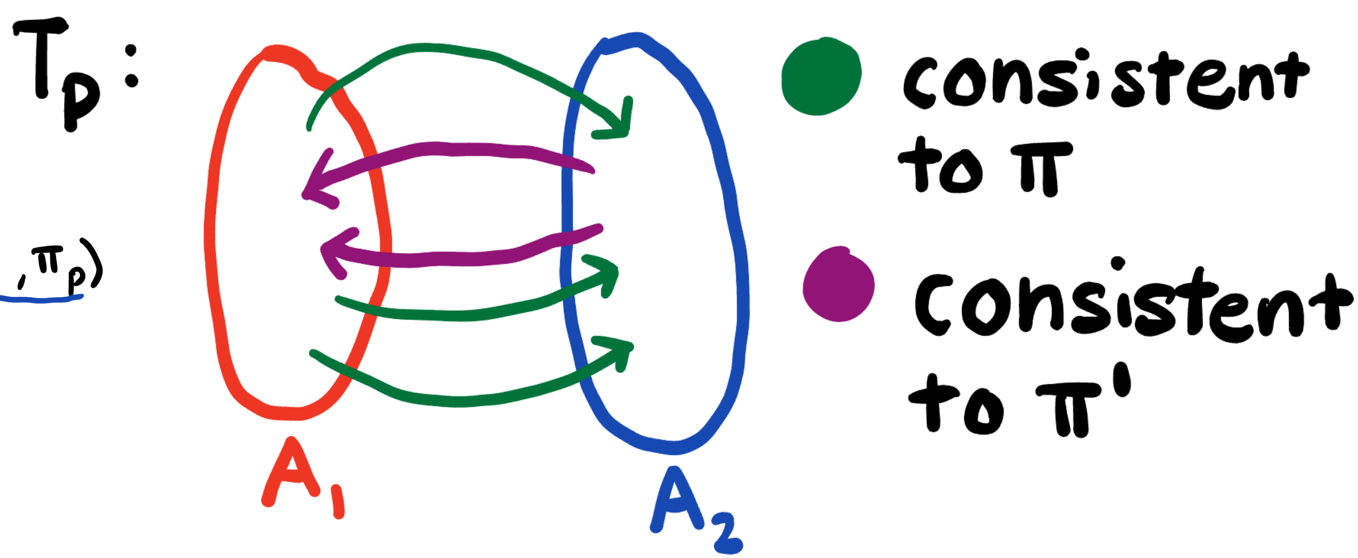

The Big Theorem: $c(T_p) \le \frac{1}{2}\binom{p}{2} + O(p^{3/2}\log p)$

If r is the smallest integer with $2^r \ge p$, since A_1, A_2 partition the vertex set in half, we have $|A_1|, |A_2| \le 2^{r-1}$. Using **Lemma 2**:

$$e(A_1, A_2) - e(A_2, A_1) \le |A_1|^{1/2}|A_2|^{1/2}p^{1/2}$$
$$\boxed{\le 2^{r-1}p^{1/2}}$$

The Big Theorem: $c(T_p) \leq \frac{1}{2}\binom{p}{2} + O(p^{3/2}\log p)$

Now repeat the process on A_1 and A_2. Let A_{11}, A_{12} partition A_1 in half. And correspondingly for A_{21}, A_{22}.

Again, $e(A_{11}, A_{12})$ counts the edges between A_{11}, A_{12} that are consistent in π. Meanwhile $e(A_{12}, A_{11})$ are all the edges consistent in π'.

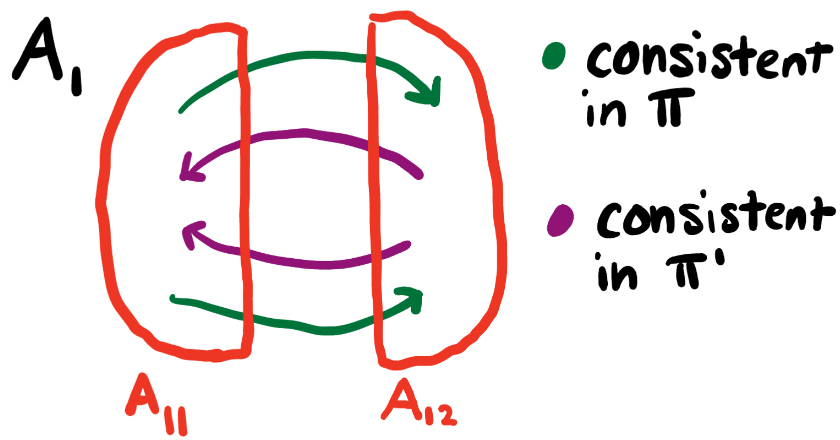

$$\boxed{\textbf{The Big Theorem}: c(T_p) \leq \frac{1}{2}\binom{p}{2} + O(p^{3/2}\log p)}$$

In general, on the kth step, we partition π into 2^k chunks, with 2^{k-1} consecutive blocks of the form $A_{\epsilon 1}, A_{\epsilon 2}$. The sum over all 2^{k-1} of these $e(A_{\epsilon 1}, A_{\epsilon 2}) - e(A_{\epsilon 2}, A_{\epsilon 1})$ is then at most

$$2^{k-1} \cdot (2^{r-k}p^{1/2}) = 2^{r-1}p^{1/2}$$

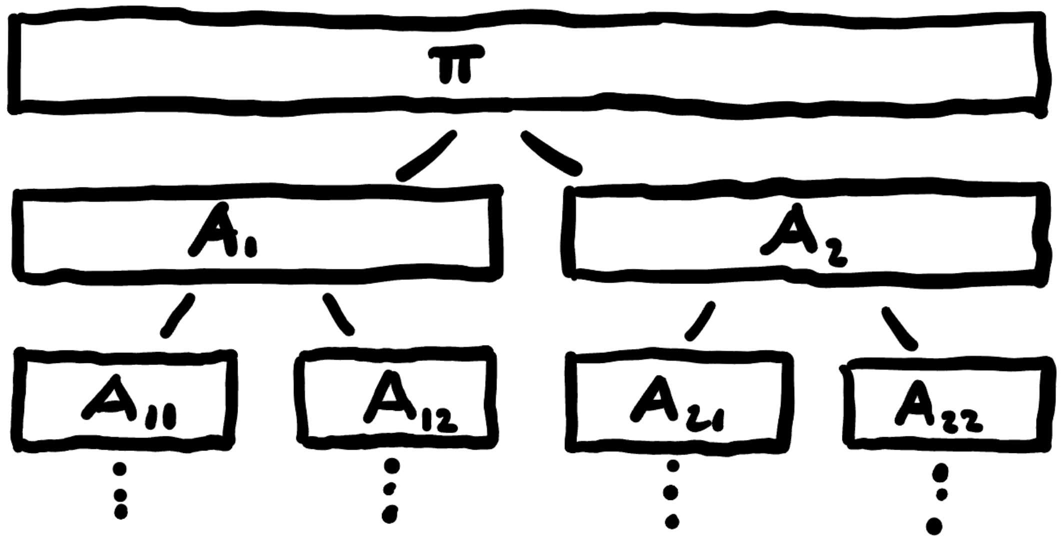

The Big Theorem: $c(T_p) \leq \frac{1}{2}\binom{p}{2} + O(p^{3/2}\log p)$

Important: Notice that each edge in T_p shows up in **precisely one** $A_{\epsilon 1}, A_{\epsilon 2}$ for some $k \in \{1, \ldots, r\}$. It's like we are binary searching for that edge! So the sum over all $k \in \{1, \ldots, r\}$ of all the $e(A_{\epsilon 1}, A_{\epsilon 2}) - e(A_{\epsilon 2}, A_{\epsilon 1})$ is precisely

$$\begin{aligned} c(\pi, T_p) - c(\pi', T_p) &\leq 2^{r-1} p^{1/2} r \\ &< p^{3/2} \log p \qquad \square \end{aligned}$$