Lecture 10: Pathsets

Instructor: Benjamin Rossman

1 Tree-width and $\max_\theta \kappa_\theta(G)$

In the last lecture notes, we showed the upper bound $\max_\theta \kappa_\theta(G) \le \mathbf{bw}(G) \le \mathbf{tw}(G) + 1$. We now show the nearly matching lower bound $\max_\theta \kappa_\theta(G) \ge \Omega(\mathbf{tw}(G)/\log \mathbf{tw}(G))$.

Definition 1. For $s, t \in V(G)$, let $\mathrm{Path}_{s,t}(G)$ be the set of directed st-paths in G. Let $\mathrm{Path}(G) := \bigcup_{s,t \in V(G)} \mathrm{Path}_{s,t}(G)$.

A *concurrent flow* on G is a non-negative function $F : \mathrm{Path}(G) \to \mathbb{R}_{\ge 0}$. For sets $S, T \subseteq V(G)$, let

$$F(S,T) := \sum_{s \in S,\, t \in T,\, \pi \in \mathrm{Path}_{s,t}(G)} F(\pi).$$

The following lemma describes an obstruction to tree-width in the form of a certain concurrent flow.

Lemma 2 (Feige-Hajiaghayi-Lee (2008) and Marx (2010))**.**
For every graph G with tree-width $\ge k$, there exists a set $W \subseteq V(G)$ with $|W| \ge \frac{2}{3}k$ and a concurrent flow F such that

- $\sum_{\pi \in \mathrm{Path}(G)\,:\, v \in V(\pi)} F(\pi) \le 1$ *for all $v \in V$ (that is, F has vertex-capacity 1),*
- *F routes $\ge \frac{1}{ck \log k}$ flow between every pair of distinct vertices in W where $c > 0$ is an absolute constant (that is, $F(\{s\},\{t\}) \ge \frac{1}{ck\log k}$ for all distinct $s, t \in W$). Note that this implies $F(S,T) \ge \frac{|S|\cdot|T|}{ck \log k}$ for all disjoint $S, T \subseteq W$.*

To get some intuition for Lemma 2, consider the special case where G has a topological clique-minor of size $k+1$ (which implies $\mathbf{tw}(G) \ge k$). If W is the set of degree-k vertices in the clique-minor, then there is a family of disjoint paths (i.e. disjoint except for endpoints) between all pairs of vertices in W. By routing $1/2k$ flow on each of these paths, we get a flow F with vertex-capacity 1 that routes $1/2k$ flow between every pair of distinct vertices in W.

It is possible to have large tree-width without having a large clique-minor. Nevertheless, Lemma 2 implies that having tree-width $\ge k$ implies the existence of a pair (W, F) that looks like a *fractional* version of a topological $\frac{2}{3}k$-clique minor.

The next theorem takes the pair (W, F) from Lemma 2 and converts it into a threshold weighting θ and a hitting set $\mathcal{H}$ that witnesses $\kappa_\theta(G) \ge \Omega(\mathbf{tw}(G)/\log \mathbf{tw}(G))$.

Theorem 3. $\max_\theta \kappa_\theta(G) \geq \Omega(\mathbf{tw}(G)/\log \mathbf{tw}(G))$

Proof. Let $k = \mathbf{tw}(G)$ and fix W, F as in Lemma 2. Consider the function $M_0 : V(G) \times V(G) \to \mathbb{R}_{\geq 0}$ defined by

$$M_0(v, w) := \sum_{\pi \in \mathrm{Path}(G)\,:\,(v,w) \in E(\pi)} F(\pi).$$

That is $M_0(v, w)$ is the amount of F-flow that routes through the directed edge (v, w). Since F has vertex-capacity 1, it follows that $\sum_{w \in V(G)\,:\,\{v,w\} \in E(G)} M_0(v, w) \leq 1$ for all $v \in V(G)$. Also note that $M_0(v, w) = 0$ whenever $\{v, w\} \notin E(G)$. By increasing the value of $M_0(v, w)$ on certain pairs (v, w) with $\{v, w\} \in E(G)$, we may obtain a Markov chain $M : V(G) \to V(G) \to [0, 1]$ on G with the property that $M(v, w) \geq M_0(v, w)$ for all v, w. (Recall that a Markov chain on G satisfies $\sum_w M(v, w) = 1$ for all v, and $M(v, w) = 0$ for all $\{v, w\} \notin E(G)$.)

For every $H \subseteq G$, we have

$$\begin{aligned}
\Delta_\theta(H) &= \sum_{(v,w)\,:\,v \in V(H),\, \{v,w\} \in E(G) \setminus E(H)} M(v, w) \quad \text{(by Lecture 7, Example 14)} \\
&\geq \sum_{(v,w)\,:\,v \in V(H),\, \{v,w\} \in E(G) \setminus E(H)} \ \sum_{\pi \in \mathrm{Path}(G)\,:\,(v,w) \in E(\pi)} F(\pi) \quad \text{(since } M(v, w) \geq M_0(v, w)) \\
&\geq \sum_{(s,t) \in V(H) \times (V(G) \setminus V(H))} \ \sum_{\pi \in \mathrm{Path}_{s,t}(G)} F(\pi) \\
&= F(V(H), V(G) \setminus V(H)).
\end{aligned}$$

To see why this last inequality holds, for each $\pi \in \mathrm{Path}_{s,t}(G)$, we "charge" π to the first directed edge $(v, w) \in E(\pi)$ with $v \in V(H)$ and $\{v, w\} \in E(G) \setminus E(H)$.

Let $\mathcal{H}$ be the set of subgraphs $H \subseteq G$ such that $\frac{2}{9}k \leq |W \cap V(H)| < \frac{4}{9}k$. Note that $\mathcal{H}$ is a hitting set for $\mathrm{Pattern}(G)$. Via the dual expression for $\kappa_\theta(G)$, we have

$$\begin{aligned}
\kappa_\theta(G) \geq \min_{H \in \mathcal{H}} \Delta_\theta(H) \geq \min_{H \in \mathcal{H}} F(V(H), V(G) \setminus V(H)) &\geq \min_{H \in \mathcal{H}} \frac{|W \cap V(H)| \cdot |W \setminus V(H)|}{ck \log k} \\
&\geq \min_{\frac{2}{9}k \leq \ell \leq \frac{4}{9}k} \frac{\ell \cdot (\frac{2}{3}k - \ell)}{ck \log k} \\
&= \frac{8}{81c} \cdot \frac{k}{\log k}.
\end{aligned}$$

□

We remark that, in special cases such as when $G = K_k$, or when G is a constant-degree expander, we get a stronger bound $\max_\theta \kappa_\theta(G) = \Omega(\mathbf{tw}(G))$ (without the $\log \mathbf{tw}(G)$-factor loss).

Conjecture 4. $\max_\theta \kappa_\theta(G) = \Omega(\mathbf{tw}(G))$ *for all graphs* G.

Conjecture 4 is a purely graph-theoretic question. (I am happy to think about this question with anyone who's interested.)

2 Corollary: Size Hierarchy Theorem for AC^0

A special case of our main result:

Theorem 5. *The average-case* AC^0 *complexity of* $\mathrm{SUB}(K_k)$ *on* $\mathbf{X}_{2/(k-1)}$ *is* $\Omega(n^{k/4})$ *and* $O(n^{k/4+2})$. *(Moreover, the upper bound is uniform.)*

We get identical bounds for k-CLIQUE (a.k.a. $\mathrm{SUB}_{\mathrm{uncolored}}(K_k)$) on the Erdos-Renyi random graph $\mathbf{G}(n, n^{-2/(k-1)})$. (See discussion in Lecture 7.)

Corollary 6. *The same bounds apply to the average-case* AC^0 *complexity of* k-CLIQUE *on* $\mathbf{G}(n, n^{-2/(k-1)})$.

Our lower bound for $\mathrm{SUB}(K_k)$ immediately implies a "Size Hierarchy Theorem for Uniform AC^0": for every k,

$$\text{uniform-}\mathsf{AC}^0[\,\text{size } O(n^{k/4})\,] \subsetneq \text{uniform-}\mathsf{AC}^0[\,\text{size } O(n^{k})\,].$$

An even stronger "incompressibility" result that follows from our upper and lower bounds is the following:

Corollary 7. *For every* $\varepsilon > 0$, *there exists a Boolean function (a padded version of average-case* $O(1/\varepsilon)$-CLIQUE*) that is computable by uniform* AC^0 *circuits of depth* $O(1/\varepsilon)$ *and size* $O(n^{1+\varepsilon})$, *but not computable by non-uniform* AC^0 *circuits of depth* $o(\frac{\log n}{\log\log n})$ *and size* $O(n)$.

In other words, uniform AC^0 circuit size $O(n^{1+\varepsilon})$ cannot be compressed to non-uniform AC^0 circuit size $O(n)$, even allowing much greater depth. (Amano (2010) improved the depth $O(1/\varepsilon)$ in Corollary 6 to the constant 2 by extending our lower bound technique to the k-CLIQUE problem on ℓ-hypergraphs.)

3 Formulas vs. Circuits

We have seen that the parameter $\kappa_\theta(G)$ characterizes the average-case AC^0 <u>circuit</u> size of $\mathrm{SUB}(G)$? What about AC^0 <u>formula</u> size? Of particular interest is the path $G = P_k$. This graph has tree-width 1, and indeed $\mathrm{SUB}(P_k)$ is solvable (even in the worse-case) by AC^0 circuits of size $n^{O(1)}$. In fact, the "recursive doubling" technique produces AC^0 circuits of size $O(n^3)$ and depth $O(\log k)$. This converts to AC^0 formula size $n^{O(\log k)}$. (Recall that every depth d circuit of size s is equivalent to depth d circuit of size at most s^d.) Over the remaining 2.5 lectures, we will show that AC^0 formulas $\mathrm{SUB}(P_k)$ require size $n^{\Omega(\log k)}$, even up to depth $o(\frac{\log n}{\log\log n})$. This result separates the power of AC^0 circuits vs. formula in a very strong way.

Although $G = P_k$ is our main application, our lower bound technique applies to average-case $\mathrm{SUB}(G)$ on $\mathbf{X}_\theta$ for any graph G and threshold weighting θ. Similar to $\kappa_\theta(G)$, we define a parameter $\tau_\theta(G)$ and show the following:

1. the average-case AC^0 <u>formula</u> size of $\mathrm{SUB}(G)$ on $\mathbf{X}_\theta$ is at least $n^{\tau_\theta(G)}$,
2. $\max_\theta \tau_\theta(P_k) \geq \Omega(\log k)$,
3. for all graphs G, $\max_\theta \tau_\theta(G) \geq \mathbf{td}(G)^{\Omega(1)}$ where $\mathbf{td}(G)$ is the *tree-depth* of G (a graph parameter related to tree-width).

Compared to $\kappa_\theta(G)$, the parameter $\tau_\theta(G)$ is significantly more complicated to define.

In order to motivate the definition of $\tau_\theta(G)$, we will take a second look at the proof of our $n^{\kappa_\theta(G)-o(1)}$ lower bound for AC^0 circuit size. In that proof, we had two independent random objects: the graph $\mathbf{X}_\theta$ and the tuple $\boldsymbol{\alpha} \in [n]^{V(G)}$. Our argument was based on the family of events $\mathrm{ALL}(g{\restriction}R_{\mathbf{X}_\theta,H(\boldsymbol{\alpha})})$ over gates g in the circuit (after being converted to fan-in 2) and subgraphs $H \subseteq G$. For our formula size lower bound, it is not enough to consider the event $\mathrm{ALL}(g{\restriction}R_{\mathbf{X}_\theta,H(\boldsymbol{\alpha})})$ for a single random $\boldsymbol{\alpha}$. Rather, we keep track of sets $\{\alpha \in [n]^{V(H)} : \mathrm{ALL}(g{\restriction}R_{\mathbf{X}_\theta,H(\alpha)})\}$ (which we regard as a "$V(H)$-ary relations") where $\mathbf{X}_\theta$ alone is the only random object. This perspective enables us to state stronger versions of the two technical lemmas from the previous lecture.

4 Strengthening Our Main Technical Lemmas

Henceforth, we fix an arbitrary connected graph G and strict threshold weighting θ. To simplify notation, we will suppress θ by writing $\Delta(H)$ instead of $\Delta_\theta(H)$, and $\mathbf{X}$ instead of $\mathbf{X}_\theta$.

Definition 8 (The Relation $\mathcal{A}_X(f,H)$)**.**
For every function $f : \{0,1\}^{E(G^{\uparrow n})} \to \{0,1\}$ and graphs $X \subseteq G^{\uparrow n}$ and $H \subseteq G$, we define the relation $\mathcal{A}_X(f,H) \subseteq [n]^{V(H)}$ by

$$\mathcal{A}_X(f,H) := \{\alpha \in [n]^{V(H)} : \mathrm{ALL}(f{\restriction}R_{X,H(\alpha)})\}.$$

Note the density of this relation:

$$\mu(\mathcal{A}_X(f,H)) = \mathop{\mathbb{P}}_{\boldsymbol{\alpha}\in[n]^{V(H)}}[\ \mathrm{ALL}(f{\restriction}R_{X,H(\boldsymbol{\alpha})})\ ].$$

We now restate (in black) and strengthen (in red) the two technical lemmas of the last lecture.

Lemma 9 (restated & strengthened)**.**

1. *If $f : \{0,1\}^{E(G^{\uparrow n})} \to \{0,1\}$ computes $\mathrm{SUB}(G)$ a.a.s. on $\mathbf{X}$, then*

$$\mathop{\mathbb{E}}_{\mathbf{X}}\Big[\ \mu(\mathcal{A}_{\mathbf{X}}(f,G))\ \Big] \geq \frac{1}{e} - o(1), \quad \mathop{\mathbb{P}}_{\mathbf{X}}\Big[\ \mu(\mathcal{A}_{\mathbf{X}}(f,G)) \geq 0.99\ \Big] \geq \frac{1}{e} - o(1).$$

2. *If $f : \{0,1\}^{E(G^{\uparrow n})} \to \{0,1\}$ is computed by an AC^0 circuit of size $n^{O(\log\log n)}$ and depth $o(\frac{\log n}{\log\log n})$, then*

$$\mathop{\mathbb{E}}_{\mathbf{X}}\Big[\ \mu(\mathcal{A}_{\mathbf{X}}(f,H))\ \Big] \leq n^{-\Delta(H)+o(1)}, \quad \mathop{\mathbb{P}}_{\mathbf{X}}\Big[\ \mu(\mathcal{A}_{\mathbf{X}}(f,H)) \leq n^{-0.99\cdot\Delta(H)}\ \Big] \geq 1 - \exp(-n^{\Omega(1)}).$$

Statement (1) is precisely Lemma 3 of Lecture 9. In the stronger version, 0.99 may be replaced by any constant < 1. This stronger inequality follows from the proof of Lemma 3 of Lecture 9: it is a combination of observations $\mathbb{P}_{\mathbf{X}}[\,\mathrm{sub}_G(\mathbf{X}) = 0\,] \geq 1/e - o(1)$ and $\mathbb{P}_{\mathbf{X},\boldsymbol{\alpha}}[\,\mathrm{ALL}(f{\restriction}R_{\mathbf{X},G(\boldsymbol{\alpha})}) \mid \mathrm{sub}_G(\mathbf{X}) = 0\,] \geq 1 - o(1)$.

Statement (2) is precisely Lemma 1 of Lecture 9 (in the single-output setting $m = 1$). In the stronger version, 0.99 may again be replaced by any constant < 1 (obs: $(1 - o(1))\Delta(H) = \Delta(H) + o(1)$ for all $\emptyset \subset H \subset G$). The stronger inequality in red follows a similar proof; there are some additional details (involving an application of Janson's Inequality), but the general idea is similar (in particular, the Switching Lemma is used in the same way).

5 H-Pathsets

We now introduce the key notion of H-pathsets for graphs $H \subseteq G$. Roughly, an H-pathset is relation $\mathcal{A} \subseteq [n]^{V(H)}$ that satisfies a collection of density relations (one for each "closed" subgraph of H, which we define next).

Definition 10. A *closed subgraph* of H is any subgraph formed by a union of connected components of H. Notation $H_0 \subseteq_{\mathrm{cl}} H$ expresses that H_0 is a closed subgraph of H. (Obs: If H has t connected components, then it has 2^t closed subgraphs. Also note that if H is a vertex-disjoint union $H_0 \uplus H_1$, then both H_0 and H_1 are closed subgraphs of H.)

Definition 11. Henceforth (for the remainder of this lecture and the next two lecture), we fix the constant $\varepsilon := n^{-0.99}$.

Definition 12. For $H \subseteq G$, an *H-pathset* (with respect to data G, θ, ε) is a relation $\mathcal{A} \subseteq [n]^{V(H)}$ such that $\mu_{V(H_0)}(\mathcal{A}) \leq \varepsilon^{\Delta(H_0)}$ for every $H_0 \subseteq_{\mathrm{cl}} H$.

Equivalently (unpacking the definition of $\mu_{V(H_0)}(\mathcal{A})$), a relation $\mathcal{A} \subseteq [n]^{V(H)}$ is an H-pathset iff for every vertex-disjoint partition $H = H_0 \uplus H_1$ (including the case of $H_0 = H$ and $H_1 = \emptyset$), we have

$$\max_{\alpha_1 \in [n]^{V(H_1)}} \ \mathop{\mathbb{P}}_{\alpha_0 \in [n]^{V(H_0)}} [\ \alpha_0 \alpha_1 \in \mathcal{A}\ ] \leq \varepsilon^{\Delta(H_0)}.$$

(We need not consider the case where $H_0 = \emptyset$ and $H_1 = H$, since $\Delta(\emptyset) = 0$. Also, note that if $\mathcal{A}$ is an H-pathset, then so is every $\mathcal{A}' \subseteq \mathcal{A}$.)

Example 13. In the simplest case $H = G$, every relation $\mathcal{A} \subseteq [n]^{V(G)}$ is a G-pathset. This is because G is connected (by assumption) and $\Delta(G) = 0$. (Even we considered non-connected G and non-strict θ, it would still be the case that every relation $\mathcal{A} \subseteq [n]^{V(G)}$ is a G-pathset,

Example 14. In the next simplest case where H is connected, a relation $\mathcal{A} \subseteq [n]^{V(H)}$ is an H-pathset iff $\mu(\mathcal{A}) \leq \varepsilon^{\Delta(H)}$. (This is because the only nontrivial density constraint comes from the partition $H = H_0 \uplus H_1$ where $H_0 = H$ and $H_1 = \emptyset$.)

Exercise 15. The next (non-)examples illustrate the definition of H-subgraph in the setting where G is the *k-cycle graph* with vertex set $\{0, \dots, k-1\}$ and $\theta \equiv 1$. Note that every proper subgraph $H \subset G$ is a disjoint union of paths and $\Delta(H)$ is the number of disjoint paths in H. Consider the case of $H = P_{0,2} \uplus P_{5,7}$ where $P_{0,2}$ is the path on $\{0, 1, 2\}$ and $P_{5,7}$ is the path on $\{5, 6, 7\}$ (where $k > 7$). Then $\mathcal{A} \subseteq [n]^{V(H)}$ is an H-pathset iff $\mu(\mathcal{A}) \leq \varepsilon^2$ and $\mu_{\{0,1,2\}}(\mathcal{A}) \leq \varepsilon$ and $\mu_{\{5,6,7\}}(\mathcal{A}) \leq \varepsilon$.

Which of the following relations are H-pathsets? [Make sure you understand these examples!]

(a) $\{\alpha \in [n]^{V(H)} : \alpha_2 = 13 \text{ and } \alpha_6 = 19\}$ *yes*

(b) $\{\alpha \in [n]^{V(H)} : \alpha_0 = \alpha_1 \text{ and } \alpha_6 = 19\}$ *yes*

(c) $\{\alpha \in [n]^{V(H)} : \alpha_0 = \alpha_5 \text{ and } \alpha_6 = 19\}$ *yes*

(d) $\{\alpha \in [n]^{V(H)} : \alpha_0 = \alpha_1 = \alpha_2\}$ *no:* $\mu_{\{5,6,7\}}(\mathcal{A}) = 1 > \varepsilon$

(e) $\{\alpha \in [n]^{V(H)} : \alpha_0 = \alpha_1 = \alpha_7\}$ *yes*

(f) $\{\alpha \in [n]^{V(H)} : \alpha_1 = \alpha_7\}$ *no:* $\mu(\mathcal{A}) = n^{-1} > \varepsilon^2$

Lemma 9(2) states that if f is AC^0-computable, then $\mu(\mathcal{A}_{\mathbf{X}}(f,H)) \leq \varepsilon^{\Delta(H)}$ with very high probability (namely, $1 - \exp(-n^{\Omega(1)})$). In other words, $\mathcal{A}_{\mathbf{X}}(f,H)$ satisfies the top-level density constraint in the definition of H-pathset. Our final strengthening of Lemma 9(2) is the following

Lemma 16 (**Main Technical Lemma, Final Form**). *If $f : \{0,1\}^{E(G^{\uparrow n})} \to \{0,1\}$ is computed by an AC^0 circuit of size $n^{O(\log\log n)}$ and depth $o(\frac{\log n}{\log\log n})$, then*

$$\mathop{\mathbb{P}}_{\mathbf{X}}\Big[\ \mathcal{A}_{\mathbf{X}}(f,H) \text{ is an } H\text{-pathset}\ \Big] \geq 1 - \exp(-n^{\Omega(1)}).$$

Lemma 16 follows from Lemma 9(2) by a simple argument, which we leave as an exercise.

Remark 17. Unpacking definitions, we have:

$$\mathcal{A}_{\mathbf{X}}(f,H) \text{ is an } H\text{-pathset} \iff \text{for every vertex-disjoint partition } H = H_0 \uplus H_1 \text{ and } \alpha_1 \in [n]^{V(H_1)},$$
$$\mathop{\mathbb{P}}_{\alpha_0 \in [n]^{V(H_0)}}\Big[\ \mathrm{ALL}(f{\restriction}R_{\mathbf{X},H^{(\alpha_0\alpha_1)}})\ \Big] \leq \varepsilon^{\Delta(H_0)}.$$

6 The Upshot: From AC^0 Formulas to "Pathset Formulas"

We state the upshot of today's lecture:

Lemma 18. *Let F be a DeMorgan formula (fan-in 2) with size $n^{O(\log\log n)}$ and $o(\frac{\log n}{\log\log n})$ AND-OR alternations and assume that F computes $\mathrm{SUB}(G)$ a.a.s. on $\mathbf{X}_\theta$. Then there exists a graph $X \subseteq G^{\uparrow n}$ with the following four properties:*

1. *$\mu(\mathcal{A}_X(F,G)) \geq .99$,*
2. *$\mathcal{A}_X(f,H)$ is an H-pathset for every sub-formula f of F and every $H \subseteq G$,*
3. *if f is an input to F (i.e. variable or negated variable), then $\mathcal{A}_X(f,H) = \emptyset$ for every $H \subseteq G$ with $|E(H)| \geq 2$,*
4. *If f is $f_1 \wedge f_2$ or $f_1 \vee f_2$, then $\mathcal{A}_X(f,H) \subseteq \displaystyle\bigcup_{H_1,H_2\,:\,H=H_1\cup H_2} \mathcal{A}_X(f_1,H_1) \bowtie \mathcal{A}_X(f_2,H_2)$.*

Proof. Conditions (3) and (4) hold for every graph $X \subseteq G^{\uparrow n}$ whatsoever (we justify this claim in a moment). Condition (2) hold for random $\mathbf{X}$ with probability $1 - o(1)$ (taking a union bound—over the $n^{O(\log\log n)}$ many sub-formulas f and $O(1)$ subgraphs $H \subseteq H$—of the "bad" event that $\mathcal{A}_{\mathbf{X}}(f,H)$ is not an H-pathset, which occurs with probability $\exp(-n^{\Omega(1)})$ by Lemma 16). Condition (1) holds for random $\mathbf{X}$ with probability $\geq \frac{1}{e} - o(1)$. Therefore, by the probabilistic method, there exists $X \subseteq G^{\uparrow n}$ satisfying (1)–(4).

Justifying condition (3): if f is an input to F corresponding to a [negated] indicator variable for an edge $\{v^{(i)}, w^{(j)}\}$ of $G^{\uparrow n}$, then $\mathrm{ALL}(f{\restriction}R_{X,H^\alpha})$ can only hold when $H = \emptyset$ or H is the single-edge graph with $E(H) = \{\{v,w\}\}$. Justifying condition (4): suppose f is $f_1 \wedge f_2$ or $f_1 \vee f_2$ and consider any $\alpha \in \mathcal{A}_X(f,H)$, that is, $\alpha \in [n]^{V(H)}$ with $\mathrm{ALL}(f{\restriction}R_{X,H^\alpha})$. Let $H_i \subseteq H$ be the graph with $E(H_i)$ being the set of $\{v,w\} \in E(H)$ such that $f_i{\restriction}R_{X,H^\alpha} : \{0,1\}^{E(H^{(\alpha)})}$ depends on the coordinate $\{v^{(\alpha_v)}, w^{(\alpha_w)}\}$. Then we have $H = H_1 \cup H_2$. Moreover, for $i = 1,2$, we have $\mathrm{ALL}(f{\restriction}R_{X,(H_i)^{\alpha_{V(H_i)}}})$, that is, $\alpha_{V(H_i)} \in \mathcal{A}_X(f_i,H_i)$. By definition of $\bowtie$ (see handout), this shows that $\alpha \in \mathcal{A}_X(f_1,H_1) \bowtie$

$\mathcal{A}_X(f_2, H_2)$. It follows that $\mathcal{A}_X(f, H) \subseteq \bigcup_{H_1, H_2 \,:\, H = H_1 \cup H_2} \mathcal{A}_X(f_1, H_1) \bowtie \mathcal{A}_X(f_2, H_2)$. (Obs: H_1, H_2 depend on the choice of $\alpha \in \mathcal{A}_X(f, H)$, so we need to take the union over all pairs H_1, H_2 with $H = H_1 \cup H_2$. Note that H_1, H_2 need not be disjoint nor closed subgraphs of H; we only require that their union equals cover H.) $\square$

With Lemma 18 in hand, we are free to forget about the following notions: AC^0 formulas, the graph $\mathbf{X}$, the event ALL$(\cdot)$. The only data we need going forward: we have a binary tree F together with a family of pathsets $\{\mathcal{A}(f, H)\}_{f \in V(F),\, H \subseteq G}$ satisfying conditions (1)–(4) of Lemma 18. Based on this combinatorial data alone (which we may informally call a "pathset formula"), our goal is to prove the best possible *lower bound* on the size of the tree F (in terms of G, θ, ε, which are the parameters in the definition of H-pathset). We achieve a strong lower bound via a notion of "pathset complexity" that we will introduce in the next lecture.