

Unapproximability of TSP^a

Theorem 80 *The approximation threshold of TSP is 1 unless $P = NP$.*

- Suppose there is a polynomial-time ϵ -approximation algorithm for TSP for some $\epsilon < 1$.
- We shall construct a polynomial-time algorithm for the NP-complete HAMILTONIAN CYCLE.
- Given any graph $G = (V, E)$, construct a TSP with $|V|$ cities with distances

$$d_{ij} = \begin{cases} 1, & \text{if } \{i, j\} \in E \\ \frac{|V|}{1-\epsilon}, & \text{otherwise} \end{cases}$$

^aSahni and Gonzales (1976).

The Proof (concluded)

- Run the alleged approximation algorithm on this TSP.
- Suppose a tour of cost $|V|$ is returned.
 - This tour must be a Hamiltonian cycle.
- Suppose a tour with at least one edge of length $\frac{|V|}{1-\epsilon}$ is returned.
 - The total length of this tour is $> \frac{|V|}{1-\epsilon}$.
 - Because the algorithm is ϵ -approximate, the optimum is at least $1 - \epsilon$ times the returned tour's length.
 - The optimum tour has a cost exceeding $|V|$.
 - Hence G has no Hamiltonian cycles.

KNAPSACK Has an Approximation Threshold of Zero^a

Theorem 81 *For any ϵ , there is a polynomial-time ϵ -approximation algorithm for KNAPSACK.*

- We have n weights $w_1, w_2, \dots, w_n \in \mathbb{Z}^+$, a weight limit W , and n values $v_1, v_2, \dots, v_n \in \mathbb{Z}^+$.^b
- We must find an $S \subseteq \{1, 2, \dots, n\}$ such that $\sum_{i \in S} w_i \leq W$ and $\sum_{i \in S} v_i$ is the largest possible.

^aIbarra and Kim (1975).

^bIf the values are fractional, the result is slightly messier but the main conclusion remains correct. Contributed by Mr. Jr-Ben Tian (R92922045) on December 29, 2004.

The Proof (continued)

- Let

$$V = \max\{v_1, v_2, \dots, v_n\}.$$

- Clearly, $\sum_{i \in S} v_i \leq nV$.
- Let $0 \leq i \leq n$ and $0 \leq v \leq nV$.
- $W(i, v)$ is the minimum weight attainable by selecting some of the first i items with a total value of v .
- Set $W(0, v) = \infty$ for $v \in \{1, 2, \dots, nV\}$ and $W(i, 0) = 0$ for $i = 0, 1, \dots, n$.^a

^aContributed by Mr. Ren-Shuo Liu (D98922016) and Mr. Yen-Wei Wu (D98922013) on December 28, 2009.

The Proof (continued)

- Then, for $0 \leq i < n$,

$$W(i + 1, v) = \min\{W(i, v), W(i, v - v_{i+1}) + w_{i+1}\}.$$

- Finally, pick the largest v such that $W(n, v) \leq W$.
- The running time is $O(n^2V)$, not polynomial time.
- Key idea: Limit the number of precision bits.

The Proof (continued)

- Define

$$v'_i = 2^b \left\lfloor \frac{v_i}{2^b} \right\rfloor.$$

- This is equivalent to zeroing each v_i 's last b bits.

- From the original instance

$$x = (w_1, \dots, w_n, W, v_1, \dots, v_n),$$

define the approximate instance

$$x' = (w_1, \dots, w_n, W, v'_1, \dots, v'_n).$$

The Proof (continued)

- Solving x' takes time $O(n^2V/2^b)$.
 - The algorithm only performs subtractions on the v_i -related values.
 - So the b last bits can be *removed* from the calculations.
 - That is, use $v'_i = \lfloor \frac{v_i}{2^b} \rfloor$ in the calculations.
 - Then multiply the returned value by 2^b .
- The solution S' is close to the optimum solution S :

$$\sum_{i \in S'} v_i \geq \sum_{i \in S'} v'_i \geq \sum_{i \in S} v'_i \geq \sum_{i \in S} (v_i - 2^b) \geq \sum_{i \in S} v_i - n2^b.$$

The Proof (continued)

- Hence

$$\sum_{i \in S'} v_i \geq \sum_{i \in S} v_i - n2^b.$$

- Without loss of generality, assume $w_i \leq W$ for all i .
 - Otherwise item i is redundant.
- V is a lower bound on OPT.
 - Picking any single item with value $\leq V$ is a legitimate choice.
- The relative error from the optimum is $\leq n2^b/V$:

$$\frac{\sum_{i \in S} v_i - \sum_{i \in S'} v_i}{\sum_{i \in S} v_i} \leq \frac{\sum_{i \in S} v_i - \sum_{i \in S'} v_i}{V} \leq \frac{n2^b}{V}.$$

The Proof (concluded)

- Suppose we pick $b = \lfloor \log_2 \frac{\epsilon V}{n} \rfloor$.
- The algorithm becomes ϵ -approximate (see Eq. (10) on p. 639).
- The running time is then $O(n^2 V / 2^b) = O(n^3 / \epsilon)$, a polynomial in n and $1/\epsilon$.^a

^aIt hence depends on the *value* of $1/\epsilon$. Thanks to a lively class discussion on December 20, 2006. If we fix ϵ and let the problem size increase, then the complexity is cubic. Contributed by Mr. Ren-Shan Luoh (D97922014) on December 23, 2008.

Pseudo-Polynomial-Time Algorithms

- Consider problems with inputs that consist of a collection of integer parameters (TSP, KNAPSACK, etc.).
- An algorithm for such a problem whose running time is a polynomial of the input length and the *value* (not length) of the largest integer parameter is a **pseudo-polynomial-time algorithm**.^a
- On p. 665, we presented a pseudo-polynomial-time algorithm for KNAPSACK that runs in time $O(n^2V)$.
- How about TSP (D), another NP-complete problem?

^aGarey and Johnson (1978).

No Pseudo-Polynomial-Time Algorithms for TSP (D)

- By definition, a pseudo-polynomial-time algorithm becomes polynomial-time if each integer parameter is limited to having a *value* polynomial in the input length.
- Corollary 43 (p. 344) showed that HAMILTONIAN PATH is reducible to TSP (D) with weights 1 and 2.
- As HAMILTONIAN PATH is NP-complete, TSP (D) cannot have pseudo-polynomial-time algorithms unless $P = NP$.
- TSP (D) is said to be **strongly NP-hard**.
- Many weighted versions of NP-complete problems are strongly NP-hard.

Polynomial-Time Approximation Scheme

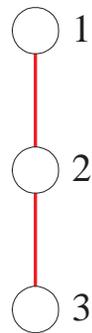
- Algorithm M is a **polynomial-time approximation scheme (PTAS)** for a problem if:
 - For each $\epsilon > 0$ and instance x of the problem, M runs in time polynomial (depending on ϵ) in $|x|$.
 - * Think of ϵ as a constant.
 - M is an ϵ -approximation algorithm for every $\epsilon > 0$.

Fully Polynomial-Time Approximation Scheme

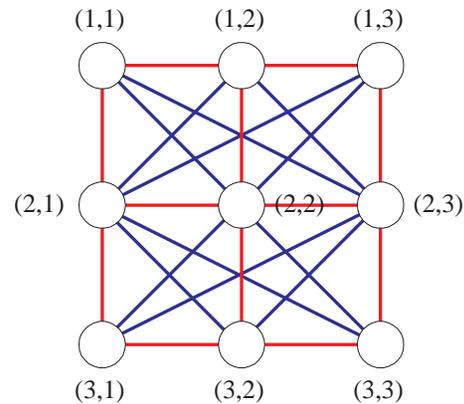
- A polynomial-time approximation scheme is **fully polynomial (FPTAS)** if the running time depends polynomially on $|x|$ and $1/\epsilon$.
 - Maybe the best result for a “hard” problem.
 - For instance, KNAPSACK is fully polynomial with a running time of $O(n^3/\epsilon)$ (p. 663).

Square of G

- Let $G = (V, E)$ be an undirected graph.
- G^2 has nodes $\{(v_1, v_2) : v_1, v_2 \in V\}$ and edges $\{\{ (u, u'), (v, v') \} : (u = v \wedge \{u', v'\} \in E) \vee \{u, v\} \in E\}$.



G

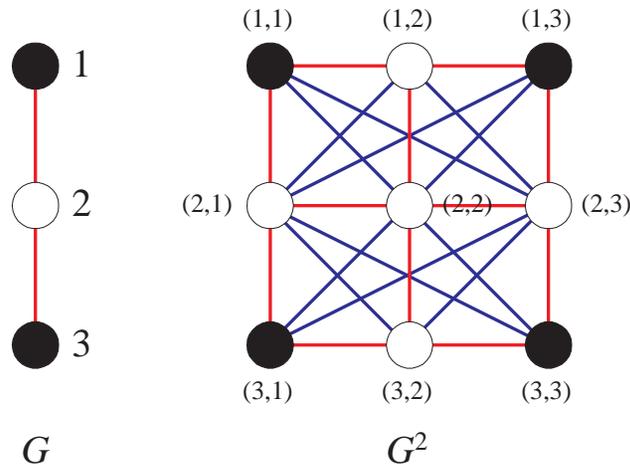


G^2

Independent Sets of G and G^2

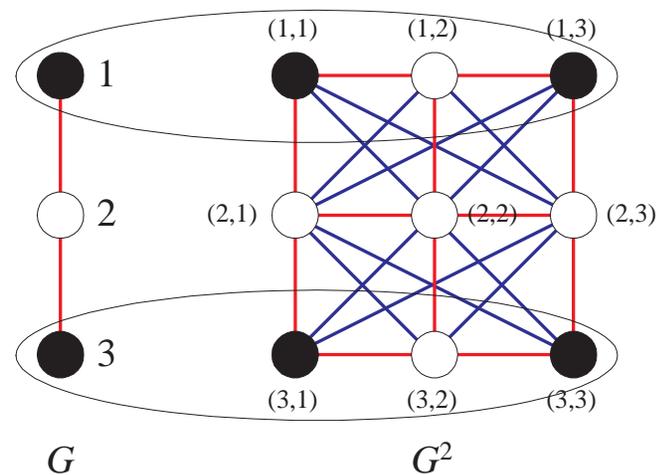
Lemma 82 $G(V, E)$ has an independent set of size k if and only if G^2 has an independent set of size k^2 .

- Suppose G has an independent set $I \subseteq V$ of size k .
- $\{(u, v) : u, v \in I\}$ is an independent set of size k^2 of G^2 .



The Proof (continued)

- Suppose G^2 has an independent set I^2 of size k^2 .
- $U \equiv \{u : \exists v \in V (u, v) \in I^2\}$ is an independent set of G .



- $|U|$ is the number of “rows” that the nodes in I^2 occupy.

The Proof (concluded)^a

- If $|U| \geq k$, then we are done.
- Now assume $|U| < k$.
- As the k^2 nodes in I^2 cover fewer than k “rows,” there must be a “row” in possession of $> k$ nodes of I^2 .
- Those $> k$ nodes will be independent in G as each “row” is a copy of G .

^aThanks to a lively class discussion on December 29, 2004.

Approximability of INDEPENDENT SET

- The approximation threshold of the maximum independent set is either zero or one (it is one!).

Theorem 83 *If there is a polynomial-time ϵ -approximation algorithm for INDEPENDENT SET for any $0 < \epsilon < 1$, then there is a polynomial-time approximation scheme.*

- Let G be a graph with a maximum independent set of size k .
- Suppose there is an $O(n^i)$ -time ϵ -approximation algorithm for INDEPENDENT SET.
- We seek a polynomial-time ϵ' -approximation algorithm with $\epsilon' < \epsilon$.

The Proof (continued)

- By Lemma 82 (p. 675), the maximum independent set of G^2 has size k^2 .
- Apply the algorithm to G^2 .
- The running time is $O(n^{2i})$.
- The resulting independent set has size $\geq (1 - \epsilon) k^2$.
- By the construction in Lemma 82 (p. 675), we can obtain an independent set of size $\geq \sqrt{(1 - \epsilon) k^2}$ for G .
- Hence there is a $(1 - \sqrt{1 - \epsilon})$ -approximation algorithm for INDEPENDENT SET by Eq. (11) on p. 640.

The Proof (concluded)

- In general, we can apply the algorithm to G^{2^ℓ} to obtain an $(1 - (1 - \epsilon)^{2^{-\ell}})$ -approximation algorithm for INDEPENDENT SET.
- The running time is $n^{2^\ell i}$.^a
- Now pick $\ell = \lceil \log \frac{\log(1-\epsilon)}{\log(1-\epsilon')} \rceil$.
- The running time becomes $n^{i \frac{\log(1-\epsilon)}{\log(1-\epsilon')}}$.
- It is an ϵ' -approximation algorithm for INDEPENDENT SET.

^aIt is not fully polynomial.

Comments

- INDEPENDENT SET and NODE COVER are reducible to each other (Corollary 40, p. 309).
- NODE COVER has an approximation threshold at most 0.5 (p. 645).
- But INDEPENDENT SET is unapproximable (see the textbook).
- INDEPENDENT SET limited to graphs with degree $\leq k$ is called k -DEGREE INDEPENDENT SET.
- k -DEGREE INDEPENDENT SET is approximable (see the textbook).

On P vs. NP

Density^a

The **density** of language $L \subseteq \Sigma^*$ is defined as

$$\text{dens}_L(n) = |\{x \in L : |x| \leq n\}|.$$

- If $L = \{0, 1\}^*$, then $\text{dens}_L(n) = 2^{n+1} - 1$.
- So the density function grows at most exponentially.
- For a unary language $L \subseteq \{0\}^*$,

$$\text{dens}_L(n) \leq n + 1.$$

– Because $L \subseteq \{\epsilon, 0, 00, \dots, \overbrace{00 \cdots 0}^n, \dots\}$.

^aBerman and Hartmanis (1977).

Sparsity

- **Sparse languages** are languages with polynomially bounded density functions.
- **Dense languages** are languages with superpolynomial density functions.

Self-Reducibility for SAT

- An algorithm exhibits **self-reducibility** if it finds a certificate by exploiting algorithms for the *decision* version of the same problem.
- Let ϕ be a boolean expression in n variables x_1, x_2, \dots, x_n .
- $t \in \{0, 1\}^j$ is a **partial** truth assignment for x_1, x_2, \dots, x_j .
- $\phi[t]$ denotes the expression after substituting the truth values of t for $x_1, x_2, \dots, x_{|t|}$ in ϕ .

An Algorithm for SAT with Self-Reduction

We call the algorithm below with empty t .

```
1: if  $|t| = n$  then  
2:   return  $\phi[t]$ ;  
3: else  
4:   return  $\phi[t_0] \vee \phi[t_1]$ ;  
5: end if
```

The above algorithm runs in exponential time, by visiting all the partial assignments (or nodes on a depth- n binary tree).

NP-Completeness and Density^a

Theorem 84 *If a unary language $U \subseteq \{0\}^*$ is NP-complete, then $P = NP$.*

- Suppose there is a reduction R from SAT to U .
- We use R to find a truth assignment that satisfies boolean expression ϕ with n variables if it is satisfiable.
- Specifically, we use R to prune the exponential-time exhaustive search on p. 686.
- The trick is to keep the already discovered results $\phi[t]$ in a table H .

^aBerman (1978).

```
1: if  $|t| = n$  then
2:   return  $\phi[t]$ ;
3: else
4:   if  $(R(\phi[t]), v)$  is in table  $H$  then
5:     return  $v$ ;
6:   else
7:     if  $\phi[t_0] = \text{“satisfiable”}$  or  $\phi[t_1] = \text{“satisfiable”}$  then
8:       Insert  $(R(\phi[t]), \text{“satisfiable”})$  into  $H$ ;
9:       return  $\text{“satisfiable”}$ ;
10:    else
11:      Insert  $(R(\phi[t]), \text{“unsatisfiable”})$  into  $H$ ;
12:      return  $\text{“unsatisfiable”}$ ;
13:    end if
14:  end if
15: end if
```

The Proof (continued)

- Since R is a reduction, $R(\phi[t]) = R(\phi[t'])$ implies that $\phi[t]$ and $\phi[t']$ must be both satisfiable or unsatisfiable.
- $R(\phi[t])$ has polynomial length $\leq p(n)$ because R runs in log space.
- As R maps to unary numbers, there are only polynomially many $p(n)$ values of $R(\phi[t])$.
- How many nodes of the complete binary tree (of invocations/truth assignments) need to be visited?
- If that number is a polynomial, the overall algorithm runs in polynomial time and we are done.

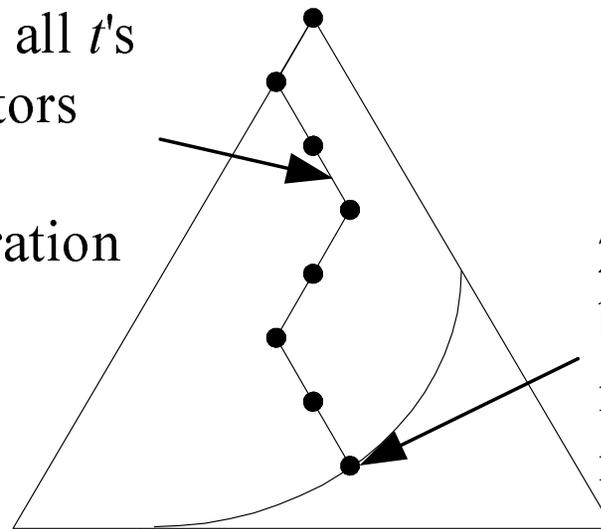
The Proof (continued)

- A search of the table takes time $O(p(n))$ in the random access memory model.
- The running time is $O(Mp(n))$, where M is the total number of invocations of the algorithm.
- The invocations of the algorithm form a binary tree of depth at most n .

The Proof (continued)

- There is a set $T = \{t_1, t_2, \dots\}$ of invocations (partial truth assignments, i.e.) such that:
 1. $|T| \geq (M - 1)/(2n)$.
 2. All invocations in T are recursive (nonleaves).
 3. None of the elements of T is a prefix of another.

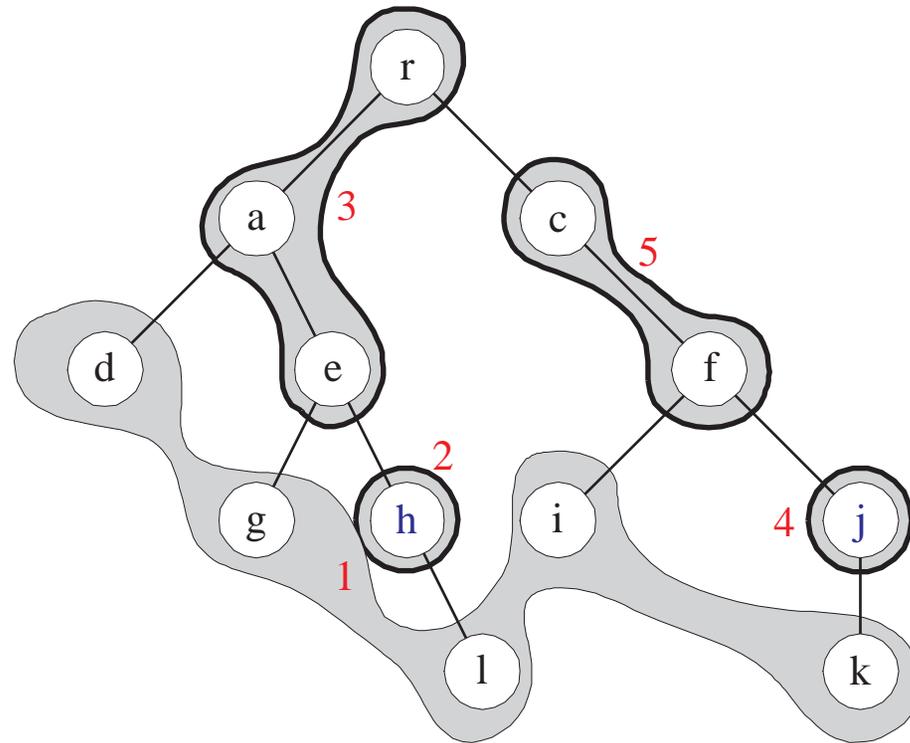
3rd step: Delete all t 's
at most n ancestors
(prefixes) from
further consideration



2nd step: Select any
bottom undeleted
invocation t and add
it to T

1st step: Delete
leaves; $(M - 1)/2$
nonleaves remaining

An Example



$$T = \{h, j\}.$$

The Proof (continued)

- All invocations $t \in T$ have different $R(\phi[t])$ values.
 - None of $h, j \in T$ is a prefix of the other.
 - The invocation of one started after the invocation of the other had terminated.
 - If they had the same value, the one that was invoked second would have looked it up, and therefore would not be recursive, a contradiction.
- The existence of T implies that there are at least $(M - 1)/(2n)$ different $R(\phi[t])$ values in the table.

The Proof (concluded)

- We already know that there are at most $p(n)$ such values.
- Hence $(M - 1)/(2n) \leq p(n)$.
- Thus $M \leq 2np(n) + 1$.
- The running time is therefore $O(Mp(n)) = O(np^2(n))$.
- We comment that this theorem holds for any sparse language, not just unary ones.^a

^aMahaney (1980).

coNP-Completeness and Density

Theorem 85 (Fortung (1979)) *If a unary language $U \subseteq \{0\}^*$ is coNP-complete, then $P = NP$.*

- Suppose there is a reduction R from SAT COMPLEMENT to U .
- The rest of the proof is basically identical except that, now, we want to make sure a formula is unsatisfiable.

Exponential Circuit Complexity

- Almost all boolean functions require

$$\frac{2^n}{2n}$$

gates to compute (generalized Theorem 14 on p. 164).

- Progress of using circuit complexity to prove exponential lower bounds for NP-complete problems has been slow.

- As of January 2006, the best lower bound is

$$5n - o(n).^a$$

^aIwama and Morizumi (2002).

Exponential Circuit Complexity for NP-Complete Problems

- We shall prove exponential lower bounds for NP-complete problems using *monotone* circuits.
 - Monotone circuits are circuits without \neg gates.
- Note that this does not settle the P vs. NP problem or any of the conjectures on p. 545.

The Power of Monotone Circuits

- Monotone circuits can only compute monotone boolean functions.
- They are powerful enough to solve a P-complete problem, MONOTONE CIRCUIT VALUE (p. 257).
- There are NP-complete problems that are not monotone; they cannot be computed by monotone circuits at all.
- There are NP-complete problems that are monotone; they can be computed by monotone circuits.
 - HAMILTONIAN PATH and CLIQUE.

CLIQUE $_{n,k}$

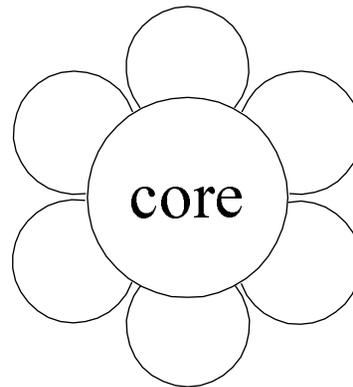
- CLIQUE $_{n,k}$ is the boolean function deciding whether a graph $G = (V, E)$ with n nodes has a clique of size k .
- The input gates are the $\binom{n}{2}$ entries of the adjacency matrix of G .
 - Gate g_{ij} is set to true if the associated undirected edge $\{i, j\}$ exists.
- CLIQUE $_{n,k}$ is a monotone function.
- Thus it can be computed by a monotone circuit.
- This does not rule out that nonmonotone circuits for CLIQUE $_{n,k}$ may use fewer gates.

Crude Circuits

- One possible circuit for $\text{CLIQUE}_{n,k}$ does the following.
 1. For each $S \subseteq V$ with $|S| = k$, there is a subcircuit with $O(k^2)$ \wedge -gates testing whether S forms a clique.
 2. We then take an OR of the outcomes of all the $\binom{n}{k}$ subsets $S_1, S_2, \dots, S_{\binom{n}{k}}$.
- This is a monotone circuit with $O(k^2 \binom{n}{k})$ gates, which is exponentially large unless k or $n - k$ is a constant.
- A **crude circuit** $\text{CC}(X_1, X_2, \dots, X_m)$ tests if *any* of $X_i \subseteq V$ forms a clique.
 - The above-mentioned circuit is $\text{CC}(S_1, S_2, \dots, S_{\binom{n}{k}})$.

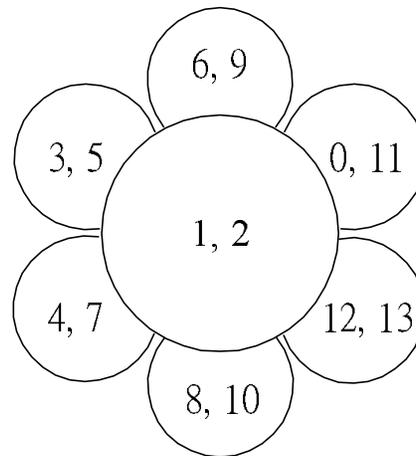
Sunflowers

- Fix $p \in \mathbb{Z}^+$ and $\ell \in \mathbb{Z}^+$.
- A **sunflower** is a family of p sets $\{P_1, P_2, \dots, P_p\}$, called **petals**, each of cardinality at most ℓ .
- All pairs of sets in the family must have the same intersection (called the **core** of the sunflower).



A Sample Sunflower

$\{\{1, 2, 3, 5\}, \{1, 2, 6, 9\}, \{0, 1, 2, 11\},$
 $\{1, 2, 12, 13\}, \{1, 2, 8, 10\}, \{1, 2, 4, 7\}\}$



The Erdős-Rado Lemma

Lemma 86 *Let \mathcal{Z} be a family of more than $M = (p - 1)^\ell \ell!$ nonempty sets, each of cardinality ℓ or less. Then \mathcal{Z} must contain a sunflower (of size p).*

- Induction on ℓ .
- For $\ell = 1$, p different singletons form a sunflower (with an empty core).
- Suppose $\ell > 1$.
- Consider a *maximal* subset $\mathcal{D} \subseteq \mathcal{Z}$ of *disjoint* sets.
 - Every set in $\mathcal{Z} - \mathcal{D}$ intersects some set in \mathcal{D} .

The Proof of the Erdős-Rado Lemma (continued)

- Suppose \mathcal{D} contains at least p sets.
 - \mathcal{D} constitutes a sunflower with an empty core.
- Suppose \mathcal{D} contains fewer than p sets.
 - Let C be the union of all sets in \mathcal{D} .
 - $|C| \leq (p-1)\ell$ and C intersects every set in \mathcal{Z} .
 - There is a $d \in C$ that intersects more than $\frac{M}{(p-1)\ell} = (p-1)^{\ell-1}(\ell-1)!$ sets in \mathcal{Z} .
 - Consider $\mathcal{Z}' = \{Z - \{d\} : Z \in \mathcal{Z}, d \in Z\}$.
 - \mathcal{Z}' has more than $M' = (p-1)^{\ell-1}(\ell-1)!$ sets.

The Proof of the Erdős-Rado Lemma (concluded)

- (continued)

- M' is just M with ℓ replaced with $\ell - 1$.

- \mathcal{Z}' contains a sunflower by induction, say

$$\{P_1, P_2, \dots, P_p\}.$$

- Now,

$$\{P_1 \cup \{d\}, P_2 \cup \{d\}, \dots, P_p \cup \{d\}\}$$

is a sunflower in \mathcal{Z} .

Comments on the Erdős-Rado Lemma

- A family of more than M sets must contain a sunflower.
- **Plucking** a sunflower entails replacing the sets in the sunflower by its core.
- By *repeatedly* finding a sunflower and plucking it, we can reduce a family with more than M sets to a family with at most M sets.
- If \mathcal{Z} is a family of sets, the above result is denoted by $\text{pluck}(\mathcal{Z})$.
- Note: $\text{pluck}(\mathcal{Z})$ is not unique.

An Example of Plucking

- Recall the sunflower on p. 703:

$$\mathcal{Z} = \{\{1, 2, 3, 5\}, \{1, 2, 6, 9\}, \{0, 1, 2, 11\}, \\ \{1, 2, 12, 13\}, \{1, 2, 8, 10\}, \{1, 2, 4, 7\}\}$$

- Then

$$\text{pluck}(\mathcal{Z}) = \{\{1, 2\}\}.$$

Razborov's Theorem

Theorem 87 (Razborov (1985)) *There is a constant c such that for large enough n , all monotone circuits for $\text{CLIQUE}_{n,k}$ with $k = n^{1/4}$ have size at least $n^{cn^{1/8}}$.*

- We shall approximate any monotone circuit for $\text{CLIQUE}_{n,k}$ by a restricted kind of crude circuit.
- The approximation will proceed in steps: one step for each gate of the monotone circuit.
- Each step introduces few errors (false positives and false negatives).
- But the resulting crude circuit has exponentially many errors.

Alexander Razborov (1963–)

