Alexander Razborov (1963–)

The Proof

- Fix $k = n^{1/4}$.
- Fix $\ell = n^{1/8}$.
- Note that[a]
$$2\binom{\ell}{2} \leq k - 1.$$
- p will be fixed later to be $n^{1/8} \log n$.
- Fix $M = (p-1)^{\ell} \ell!$.
  - Recall the Erdős-Rado lemma (p. 704).

[a]Corrected by Mr. Moustapha Bande (D98922042) on January 05, 2010.

The Proof (continued)

- Each crude circuit used in the approximation process is of the form $\mathrm{CC}(X_1, X_2, \ldots, X_m)$, where:
  - $X_i \subseteq V$.
  - $|X_i| \leq \ell$.
  - $m \leq M$.
- It answers true if any X_i is a clique.
- We shall show how to approximate any circuit for $\textsc{clique}_{n,k}$ by such a crude circuit, inductively.
- The induction basis is straightforward:
  - Input gate g_{ij} is the crude circuit $\mathrm{CC}(\{i, j\})$.

The Proof (continued)

- Any monotone circuit can be considered the OR or AND of two subcircuits.
- We shall show how to build approximators of the overall circuit from the approximators of the two subcircuits.
  - We are given two crude circuits $\mathrm{CC}(\mathcal{X})$ and $\mathrm{CC}(\mathcal{Y})$.
  - $\mathcal{X}$ and $\mathcal{Y}$ are two families of at most M sets of nodes, each set containing at most ℓ nodes.
  - We construct the approximate OR and the approximate AND of these subcircuits.
  - Then show both approximations introduce few errors.

The Proof: Positive Examples

- Error analysis will be applied to only **positive examples** and **negative examples**.
- A positive example is a graph that has $\binom{k}{2}$ edges connecting k nodes in all possible ways.
- There are $\binom{n}{k}$ such graphs.
- They all should elicit a true output from $\text{CLIQUE}_{n,k}$.

The Proof: Negative Examples

- Color the nodes with $k - 1$ different colors and join by an edge any two nodes that are colored differently.
- There are $(k-1)^n$ such graphs.
- They all should elicit a false output from $\text{CLIQUE}_{n,k}$.
  - Each set of k nodes must have 2 identically colored nodes; hence there is no edge between them.

Positive and Negative Examples with $k = 5$

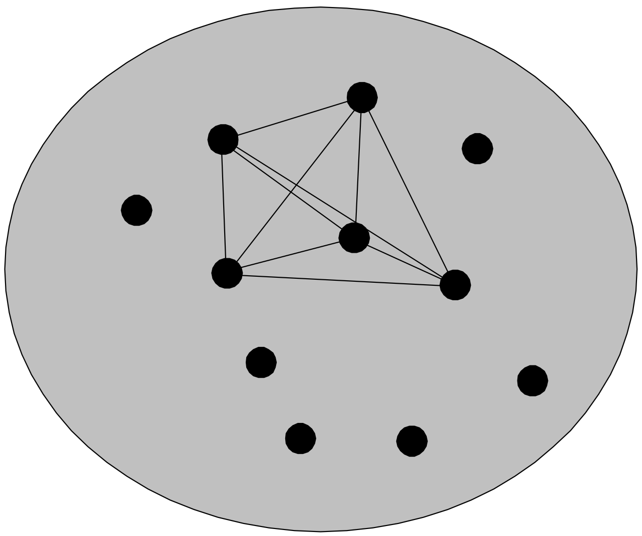

A positive example

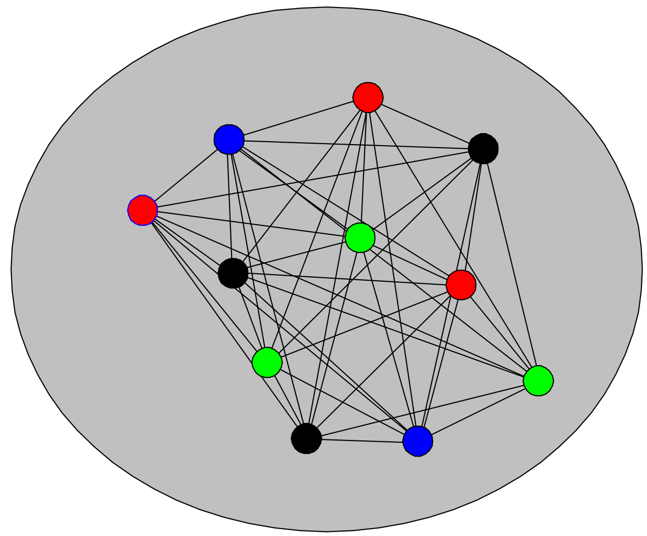

A negative example

The Proof: OR

- $CC(\mathcal{X} \cup \mathcal{Y})$ is *equivalent to* the OR of $CC(\mathcal{X})$ and $CC(\mathcal{Y})$.
- Violations occur when $|\mathcal{X} \cup \mathcal{Y}| > M$.
- Such violations can be eliminated by using

$$CC(\text{pluck}(\mathcal{X} \cup \mathcal{Y}))$$

  as the approximate OR of $CC(\mathcal{X})$ and $CC(\mathcal{Y})$.
  - Note that if $CC(\mathcal{Z})$ is true, then $CC(\text{pluck}(\mathcal{Z}))$ must be true (recall p. 702).
- We now count the number of errors this approximate OR makes on the positive and negative examples.

The Proof: OR (concluded)

- CC(pluck($\mathcal{X} \cup \mathcal{Y}$)) *introduces* a **false positive** if a negative example makes both CC($\mathcal{X}$) and CC($\mathcal{Y}$) return false but makes CC(pluck($\mathcal{X} \cup \mathcal{Y}$)) return true.
- CC(pluck($\mathcal{X} \cup \mathcal{Y}$)) *introduces* a **false negative** if a positive example makes either CC($\mathcal{X}$) or CC($\mathcal{Y}$) return true but makes CC(pluck($\mathcal{X} \cup \mathcal{Y}$)) return false.
- How many false positives and false negatives are introduced by CC(pluck($\mathcal{X} \cup \mathcal{Y}$))?

The Number of False Positives

Lemma 88 CC(pluck($\mathcal{X} \cup \mathcal{Y}$)) *introduces at most* $\frac{M}{p-1} 2^{-p}(k-1)^n$ *false positives.*

- A plucking replaces the sunflower $\{Z_1, Z_2, \ldots, Z_p\}$ with its core Z.
- A false positive is *necessarily* a coloring such that:
  - There is a pair of identically colored nodes in each petal Z_i (and so both crude circuits return false).
  - But the core contains distinctly colored nodes.
    - This implies at least one node from each same-color pair was plucked away.
- We now count the number of such colorings.

Proof of Lemma 88 (continued)

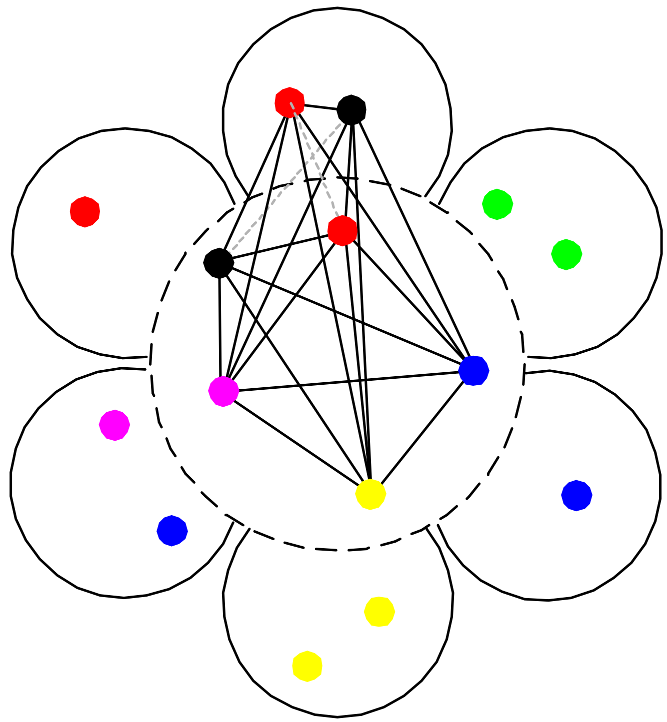

Proof of Lemma 88 (continued)

- Color nodes V at random with $k-1$ colors and let $R(X)$ denote the event that there are repeated colors in set X.
- Now $\text{prob}[R(Z_1) \wedge \cdots \wedge R(Z_p) \wedge \neg R(Z)]$ is at most

$$
\begin{aligned}
& \text{prob}[R(Z_1) \wedge \cdots \wedge R(Z_p) | \neg R(Z)] \\
= \; & \prod_{i=1}^{p} \text{prob}[R(Z_i) | \neg R(Z)] \leq \prod_{i=1}^{p} \text{prob}[R(Z_i)]. \quad (12)
\end{aligned}
$$

  - First equality holds because $R(Z_i)$ are independent given $\neg R(Z)$ as Z contains their only common nodes.
  - Last inequality holds as the likelihood of repetitions in Z_i decreases given no repetitions in $Z \subseteq Z_i$.

Proof of Lemma 88 (continued)

- Consider two nodes in Z_i.
- The probability that they have identical color is $\frac{1}{k-1}$.
- Now $\text{prob}[\, R(Z_i) \,] \leq \frac{\binom{|Z_i|}{2}}{k-1} \leq \frac{\binom{\ell}{2}}{k-1} \leq \frac{1}{2}$.
- So the probability[a] that a random coloring is a new false positive is at most 2^{-p} by inequality (12).
- As there are $(k-1)^n$ different colorings, each plucking introduces at most $2^{-p}(k-1)^n$ false positives.

---

[a] Proportion, i.e.

Proof of Lemma 88 (concluded)

- Recall that $|\mathcal{X} \cup \mathcal{Y}| \leq 2M$.
- Each plucking reduces the number of sets by $p - 1$.
- Hence at most $\frac{M}{p-1}$ pluckings occur in pluck$(\mathcal{X} \cup \mathcal{Y})$.
- At most
$$\frac{M}{p-1}\, 2^{-p} (k-1)^n$$
false positives are introduced.[a]

---

[a] Note that the numbers of errors are added not multiplied. Recall that we count how many *new* errors are introduced by each approximation step. Contributed by Mr. Ren-Shuo Liu (`D98922016`) on January 5, 2010.

The Number of False Negatives

Lemma 89 CC(pluck($\mathcal{X} \cup \mathcal{Y}$)) *introduces no false negatives.*

- Each plucking replaces a set in a crude circuit by a subset.
- This makes the test less stringent.
  - For each $Y \in \mathcal{X} \cup \mathcal{Y}$, there must exist at least one $X \in \text{pluck}(\mathcal{X} \cup \mathcal{Y})$ such that $X \subseteq Y$.
  - So if Y is a clique, then this X is also a clique.
- So plucking can only increase the number of accepted graphs.

The Number of False Negatives (concluded)

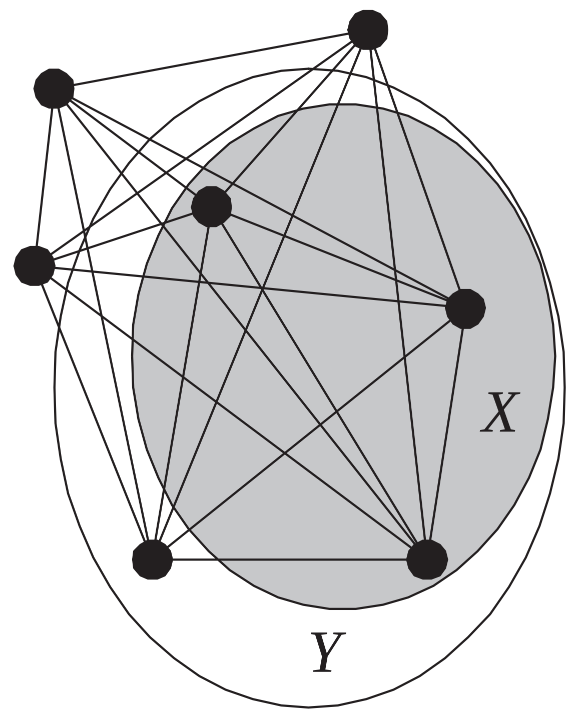

The Proof: AND

- The approximate AND of crude circuits $\text{CC}(\mathcal{X})$ and $\text{CC}(\mathcal{Y})$ is

$$\text{CC}(\text{pluck}(\{X_i \cup Y_j : X_i \in \mathcal{X}, Y_j \in \mathcal{Y}, |X_i \cup Y_j| \leq \ell\})).$$

  – Note that if $\text{CC}(\mathcal{Z})$ is true, then $\text{CC}(\text{pluck}(\mathcal{Z}))$ must be true.

- We now count the number of errors this approximate AND makes on the positive and negative examples.

The Proof: AND (concluded)

- The approximate AND *introduces* a **false positive** if a negative example makes either $\mathrm{CC}(\mathcal{X})$ or $\mathrm{CC}(\mathcal{Y})$ return false but makes the approximate AND return true.
- The approximate AND *introduces* a **false negative** if a positive example makes both $\mathrm{CC}(\mathcal{X})$ and $\mathrm{CC}(\mathcal{Y})$ return true but makes the approximate AND return false.
- How many false positives and false negatives are introduced by the approximate AND?

The Number of False Positives

Lemma 90 *The approximate* AND *introduces at most* $M^2 2^{-p}(k-1)^n$ *false positives.*

- $\text{CC}(\{X_i \cup Y_j : X_i \in \mathcal{X}, Y_j \in \mathcal{Y}\})$ introduces no false positives.
  - If $X_i \cup Y_j$ is a clique, both X_i and Y_j must be cliques, making both $\text{CC}(\mathcal{X})$ and $\text{CC}(\mathcal{Y})$ return true.
- $\text{CC}(\{X_i \cup Y_j : X_i \in \mathcal{X}, Y_j \in \mathcal{Y}, |X_i \cup Y_j| \leq \ell\})$ introduces no false positives as we are testing fewer sets for cliques.

Proof of Lemma 90 (concluded)

- $|\{X_i \cup Y_j : X_i \in \mathcal{X}, Y_j \in \mathcal{Y}, |X_i \cup Y_j| \leq \ell\}| \leq M^2$.
- Each plucking reduces the number of sets by $p - 1$.
- So pluck$(\{X_i \cup Y_j : X_i \in \mathcal{X}, Y_j \in \mathcal{Y}, |X_i \cup Y_j| \leq \ell\})$ involves $\leq M^2/(p-1)$ pluckings.
- Each plucking introduces at most $2^{-p}(k-1)^n$ false positives by the proof of Lemma 88 (p. 719).
- The desired upper bound is

$$[\,M^2/(p-1)\,]\,2^{-p}(k-1)^n \leq M^2 2^{-p}(k-1)^n.$$

The Number of False Negatives

Lemma 91 *The approximate* AND *introduces at most* $M^2\binom{n-\ell-1}{k-\ell-1}$ *false negatives.*

- We follow the same three-step proof as before.
- $\text{CC}(\{X_i \cup Y_j : X_i \in \mathcal{X}, Y_j \in \mathcal{Y}\})$ introduces no false negatives.
  - Suppose both $\text{CC}(\mathcal{X})$ and $\text{CC}(\mathcal{Y})$ accept a positive example with a clique of size k.
  - This clique must contain an $X_i \in \mathcal{X}$ and a $Y_j \in \mathcal{Y}$.
    * This is why both $\text{CC}(\mathcal{X})$ and $\text{CC}(\mathcal{Y})$ return true.
  - As the clique contains $X_i \cup Y_j$, the new circuit returns true.

Proof of Lemma 91 (continued)

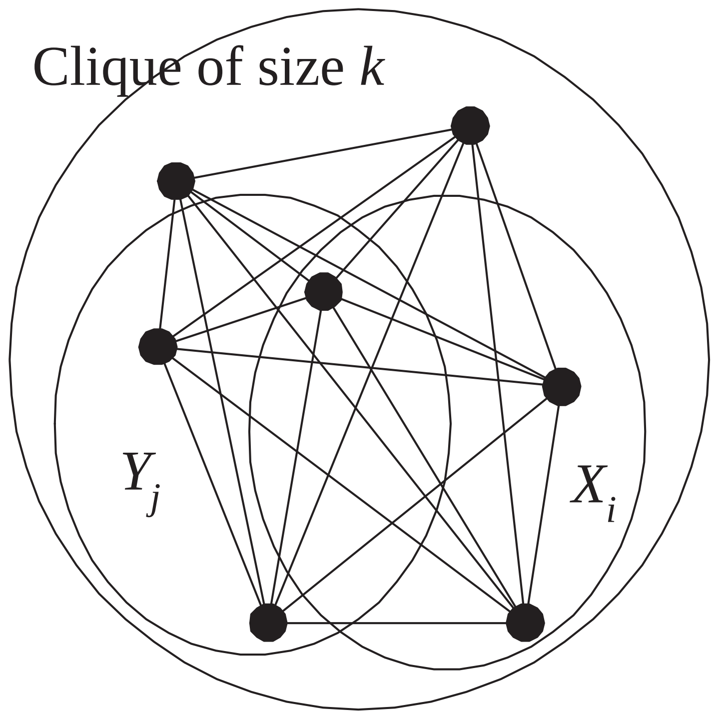

Proof of Lemma 91 (concluded)

- $\text{CC}(\{X_i \cup Y_j : X_i \in \mathcal{X}, Y_j \in \mathcal{Y}, |X_i \cup Y_j| \leq \ell\})$ introduces $\leq M^2 \binom{n-\ell-1}{k-\ell-1}$ false negatives.
  - Deletion of set $Z = X_i \cup Y_j$ larger than ℓ introduces false negatives *only if* the clique contains Z.
  - There are $\binom{n-|Z|}{k-|Z|}$ such cliques.
    * It is the number of positive examples whose clique contains Z.
  - $\binom{n-|Z|}{k-|Z|} \leq \binom{n-\ell-1}{k-\ell-1}$ as $|Z| > \ell$.
  - There are at most M^2 such Zs.
- Plucking introduces no false negatives.

Two Summarizing Lemmas

From Lemmas 88 (p. 719) and 90 (p. 728), we have:

Lemma 92 *Each approximation step introduces at most* $M^2 2^{-p}(k-1)^n$ *false positives.*

From Lemmas 89 (p. 724) and 91 (p. 730), we have:

Lemma 93 *Each approximation step introduces at most* $M^2\binom{n-\ell-1}{k-\ell-1}$ *false negatives.*

The Proof (continued)

- The above two lemmas show that each approximation step introduces "few" false positives and false negatives.
- We next show that the resulting crude circuit has "a lot" of false positives or false negatives.

The Final Crude Circuit

Lemma 94 *Every final crude circuit is:*

1. *Identically false—thus wrong on all positive examples.*
2. *Or outputs true on at least half of the negative examples.*

- Suppose it is not identically false.
- By construction, it accepts at least those graphs that have a clique on some set X of nodes, with $|X| \leq \ell$, which at $n^{1/8}$ is less than $k = n^{1/4}$.
- The proof of Lemma 88 (p. 719ff) shows that at least half of the colorings assign different colors to nodes in X.
- So half of the negative examples have a clique in X and are accepted.

The Proof (continued)

- Recall the constants on p. 711: $k = n^{1/4}$, $\ell = n^{1/8}$, $p = n^{1/8} \log n$, $M = (p-1)^\ell \ell! < n^{(1/3)n^{1/8}}$ for large n.
- Suppose the final crude circuit is identically false.
  - By Lemma 93 (p. 733), each approximation step introduces at most $M^2 \binom{n-\ell-1}{k-\ell-1}$ false negatives.
  - There are $\binom{n}{k}$ positive examples.
  - The original crude circuit for $\text{CLIQUE}_{n,k}$ has at least

$$\frac{\binom{n}{k}}{M^2 \binom{n-\ell-1}{k-\ell-1}} \geq \frac{1}{M^2} \left(\frac{n-\ell}{k} \right)^\ell \geq n^{(1/12)n^{1/8}}$$

gates for large n.

The Proof (concluded)

- Suppose the final crude circuit is not identically false.
  - Lemma 94 (p. 735) says that there are at least $(k-1)^n/2$ false positives.
  - By Lemma 92 (p. 733), each approximation step introduces at most $M^2 2^{-p}(k-1)^n$ false positives.
  - The original crude circuit for $\text{CLIQUE}_{n,k}$ has at least

$$\frac{(k-1)^n/2}{M^2 2^{-p}(k-1)^n} = \frac{2^{p-1}}{M^2} \geq n^{(1/3)n^{1/8}}$$

    gates.

P $\neq$ NP Proved?

- Razborov's theorem says that there is a monotone language in NP that has no polynomial monotone circuits.
- If we can prove that all monotone languages in P have polynomial monotone circuits, then P $\neq$ NP.
- But Razborov proved in 1985 that some monotone languages in P have no polynomial monotone circuits!

Computation That Counts

Counting Problems

- Counting problems are concerned with the number of solutions.
  - #SAT: the number of satisfying truth assignments to a boolean formula.
  - #HAMILTONIAN PATH: the number of Hamiltonian paths in a graph.
- They cannot be easier than their decision versions.
  - The decision problem has a solution if and only if the solution count is larger than 0.
- But they can be harder than their decision versions.

Decision and Counting Problems

- FP is the set of polynomial-time computable functions $f : \{0, 1\}^* \to \mathbb{Z}$.
  - GCD, LCM, matrix-matrix multiplication, etc.
- If $\#\textsc{sat} \in \text{FP}$, then $\text{P} = \text{NP}$.
  - Given boolean formula ϕ, calculate its number of satisfying truth assignments, k, in polynomial time.
  - Declare "$\phi \in \textsc{sat}$" if and only if $k \geq 1$.
- The validity of the reverse direction is open.

A Counting Problem Harder than Its Decision Version

- CYCLE asks if a directed graph contains a cycle.
- #CYCLE counts the number of cycles in a directed graph.
- CYCLE is in P by a simple greedy algorithm.
- But #CYCLE is hard unless P = NP.

Counting Class #P

A function f is in #P (or $f \in$ #P) if

- There exists a polynomial-time NTM M.
- $M(x)$ has $f(x)$ accepting paths for all inputs x.
- $f(x) =$ number of accepting paths of $M(x)$.

Some #P Problems

- $f(\phi)$ = number of satisfying truth assignments to ϕ.
  - The desired NTM guesses a truth assignment T and accepts ϕ if and only if $T \models \phi$.
  - Hence $f \in \#P$.
  - f is also called #SAT.
- #HAMILTONIAN PATH.
- #3-COLORING.

#P Completeness

- Function f is #P-complete if
  - $f \in$ #P.
  - #P $\subseteq \text{FP}^f$.
    - Every function in #P can be computed in polynomial time with access to a black box[a] or **oracle** for f.
  - Of course, oracle f will be accessed only a polynomial number of times.
  - #P is said to be **polynomial-time Turing-reducible to** f.

[a] Think of it as a subroutine

#SAT Is #P-Complete[a]

- First, it is in #P (p. 744).
- Let $f \in \#\mathrm{P}$ compute the number of accepting paths of M.
- Cook's theorem uses a *parsimonious* reduction from M on input x to an instance ϕ of SAT (p. 273).
  - Hence the number of accepting paths of $M(x)$ equals the number of satisfying truth assignments to ϕ.
- Call the oracle #SAT with ϕ to obtain the desired answer regarding $f(x)$.

[a] Valiant (1979); in fact, #2SAT is also #P-complete.

Leslie G. Valiant (1949–)

Avi Wigderson (2009), "Les Valiant singlehandedly created, or completely transformed, several fundamental research areas of computer science. [$\cdots$] We all became addicted to this remarkable throughput, and expect more."

CYCLE COVER

- A set of node-disjoint cycles that cover all nodes in a directed graph is called a **cycle cover**.

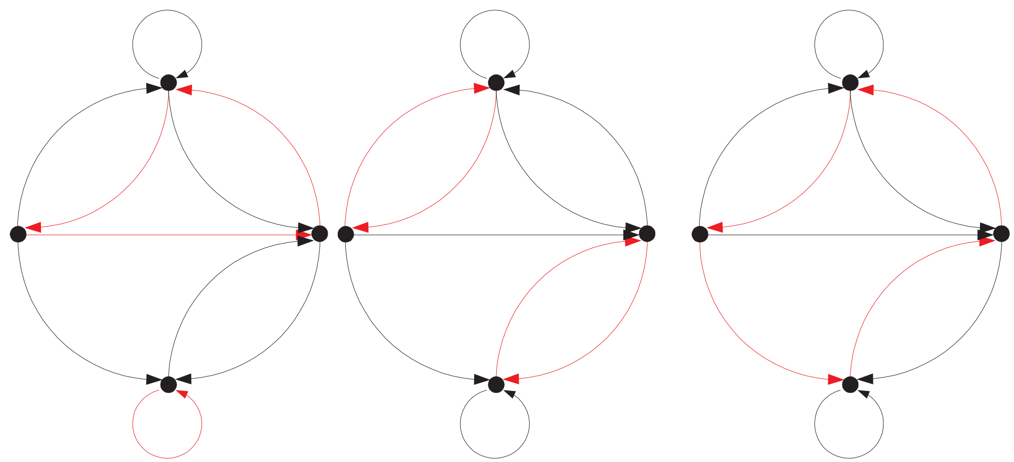

- There are 3 cycle covers (in red) above.

CYCLE COVER and BIPARTITE PERFECT MATCHING

Proposition 95 CYCLE COVER *and* BIPARTITE PERFECT MATCHING *(p. 440) are parsimoniously reducible to each other.*

- A polynomial-time algorithm creates a bipartite graph G' from any directed graph G.
- Moreover, the number cycle covers for G equals the number of bipartite perfect matchings for G'.
- And vice versa.

Corollary 96 CYCLE COVER $\in P$.

Illustration of the Proof

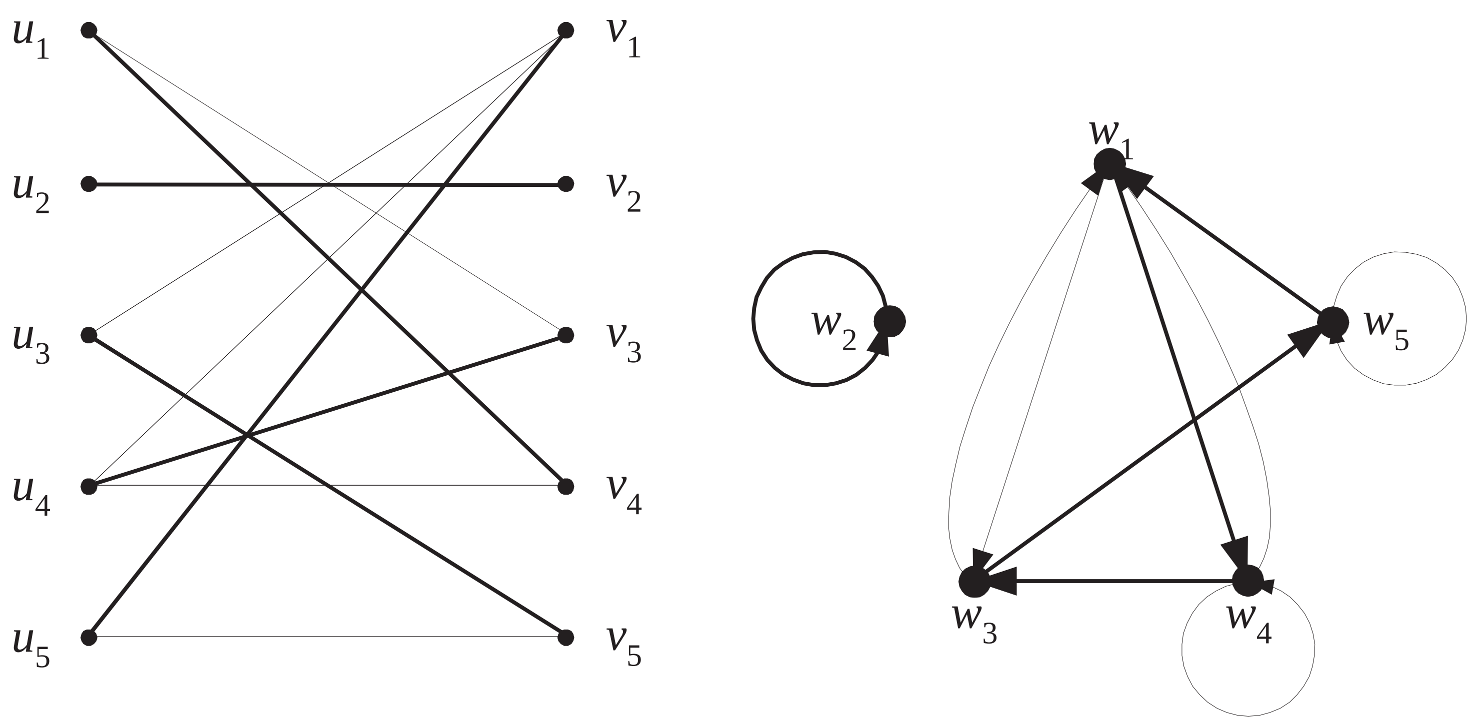

Permanent

- The **permanent** of an $n \times n$ integer matrix A is

$$\text{perm}(A) = \sum_{\pi} \prod_{i=1}^{n} A_{i,\pi(i)}.$$

  – π ranges over all permutations of n elements.

- 0/1 PERMANENT computes the permanent of a 0/1 (binary) matrix.

  – The permanent of a binary matrix is at most $n!$.

- Simpler than determinant (5) on p. 443: no signs.

- Surprisingly, much harder to compute than determinant!

Permanent and Counting Perfect Matchings

- BIPARTITE PERFECT MATCHING is related to determinant (p. 444).
- #BIPARTITE PERFECT MATCHING is related to permanent.

Proposition 97 0/1 PERMANENT *and* BIPARTITE PERFECT MATCHING *are parsimoniously reducible to each other.*

The Proof

- Given a bipartite graph G, construct an $n \times n$ binary matrix A.
  - The (i, j)th entry A_{ij} is 1 if $(i, j) \in E$ and 0 otherwise.
- Then $\text{perm}(A) =$ number of perfect matchings in G.

Illustration of the Proof Based on p. 750 (Left)

$$A = \begin{bmatrix} 0 & 0 & 1 & \boxed{1} & 0 \\ 0 & \boxed{1} & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & \boxed{1} \\ 1 & 0 & \boxed{1} & 1 & 0 \\ \boxed{1} & 0 & 0 & 0 & 1 \end{bmatrix}.$$

- $\text{perm}(A) = 4$.
- The permutation corresponding to the perfect matching on p. 750 is marked.

Permanent and Counting Cycle Covers

Proposition 98 0/1 PERMANENT *and* CYCLE COVER *are parsimoniously reducible to each other.*

- Let A be the adjacency matrix of the graph on p. 750 (right).
- Then $\text{perm}(A) =$ number of cycle covers.

Three Parsimoniously Equivalent Problems

We summarize Propositions 95 (p. 749) and 97 (p. 752) in the following.

Lemma 99 0/1 PERMANENT, BIPARTITE PERFECT MATCHING, *and* CYCLE COVER *are* **parsimoniously equivalent**.

We will show that the counting versions of all three problems are in fact #P-complete.

WEIGHTED CYCLE COVER

- Consider a directed graph G with integer weights on the edges.
- The weight of a cycle cover is the product of its edge weights.
- The **cycle count** of G is sum of the weights of all cycle covers.
  - Let A be G's adjacency matrix but $A_{ij} = w_i$ if the edge (i, j) has weight w_i.
  - Then $\text{perm}(A) = G$'s cycle count (same proof as Proposition 98 on p. 755).
- #CYCLE COVER is a special case: All weights are 1.

An Example[a]

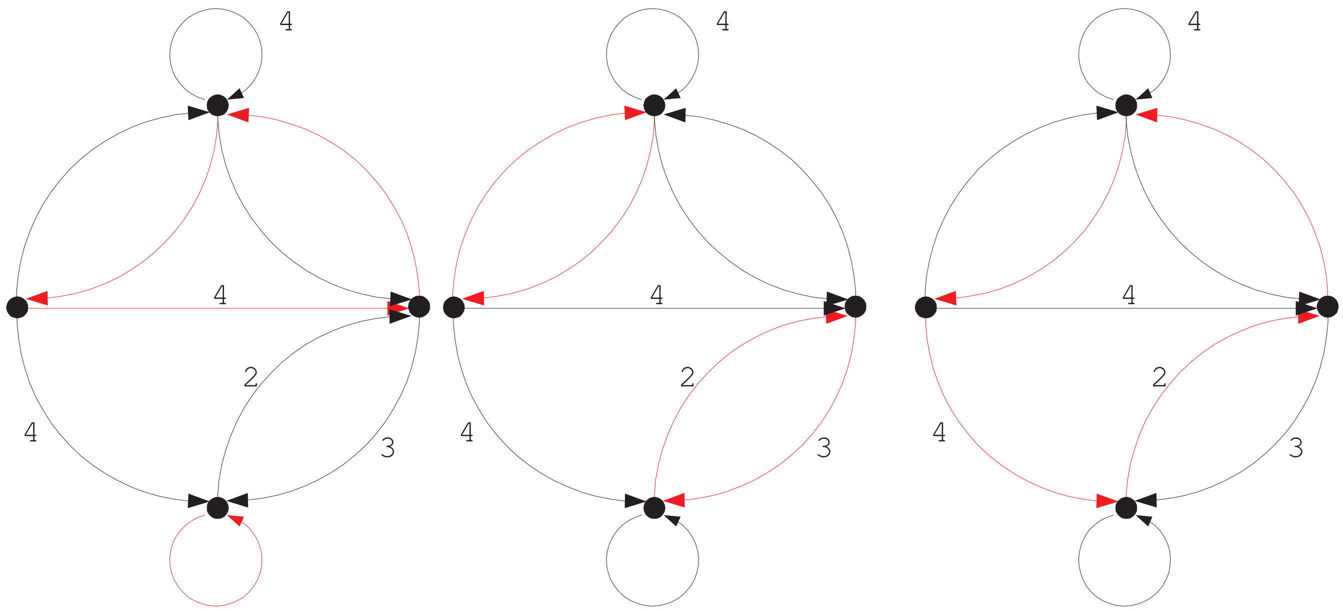

There are 3 cycle covers, and the cycle count is

$$(4 \cdot 1 \cdot 1) \cdot (1) + (1 \cdot 1) \cdot (2 \cdot 3) + (4 \cdot 2 \cdot 1 \cdot 1) = 18.$$

[a] Each edge has weight 1 unless stated otherwise.

Three #P-Complete Counting Problems

Theorem 100 (Valiant (1979)) 0/1 PERMANENT, #BIPARTITE PERFECT MATCHING, *and* #CYCLE COVER *are #P-complete.*

- By Lemma 99 (p. 756), it suffices to prove that #CYCLE COVER is #P-complete.
- #SAT is #P-complete (p. 746).
- #3SAT is #P-complete because it and #SAT are parsimoniously equivalent (p. 277).
- We shall prove that #3SAT is polynomial-time Turing-reducible to #CYCLE COVER.

The Proof (continued)

- Let ϕ be the given 3SAT formula.
  - It contains n variables and m clauses (hence $3m$ literals).
  - It has $\#\phi$ satisfying truth assignments.
- First we construct a *weighted* directed graph H with cycle count
$$\#H = 4^{3m} \times \#\phi.$$
- Then we construct an unweighted directed graph G.
- We make sure $\#H$ (hence $\#\phi$) is polynomial-time Turing-reducible to G's number of cycle covers (denoted $\#G$).

The Proof: the Clause Gadget (continued)

- Each clause is associated with a **clause gadget**.

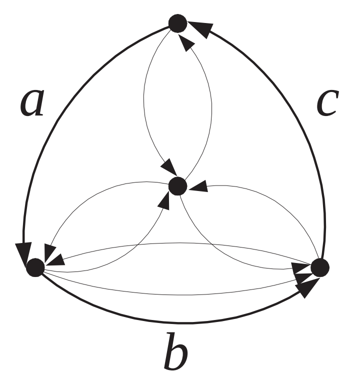

- Each edge has weight 1 unless stated otherwise.
- Each bold edge corresponds to one literal in the clause.
- There are not *parallel* lines as bold edges are schematic only (preview p. 774).

The Proof: the Clause Gadget (continued)

- Following a bold edge means making the literal false (0).
- A cycle cover cannot select *all* 3 bold edges.
  - The interior node would be missing.
- Every proper nonempty subset of bold edges corresponds to a unique cycle cover of weight 1 (see next page).

The Proof: the Clause Gadget (continued)

7 possible cycle covers, one for each satisfying assignment: (1) $a = 0, b = 0, c = 1$, (2) $a = 0, b = 1, c = 0$, etc.

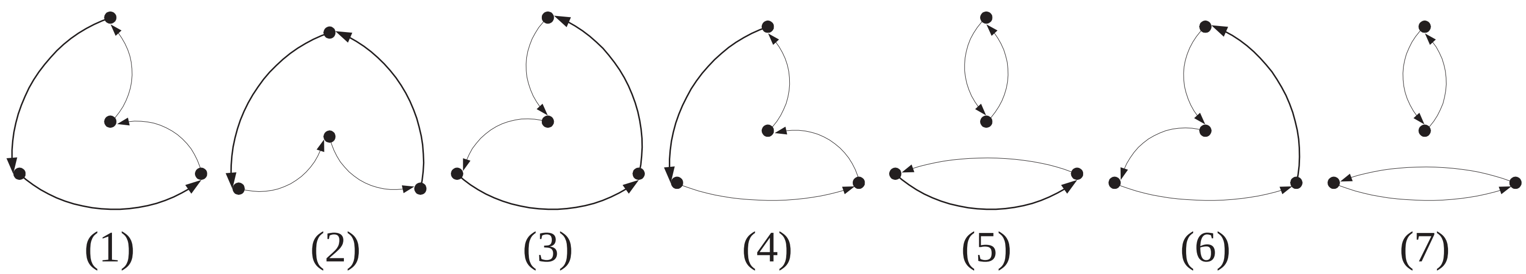

The Proof: the XOR Gadget (continued)

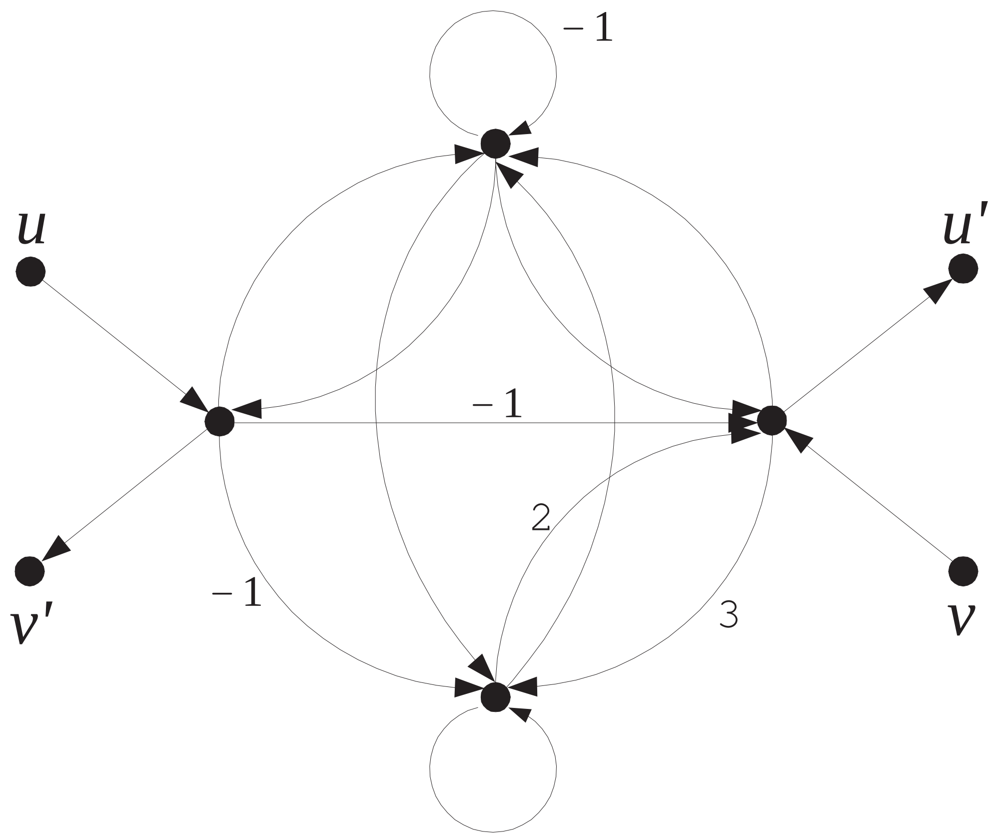

The Proof: Properties of the XOR Gadget (continued)

- The XOR gadget schema:

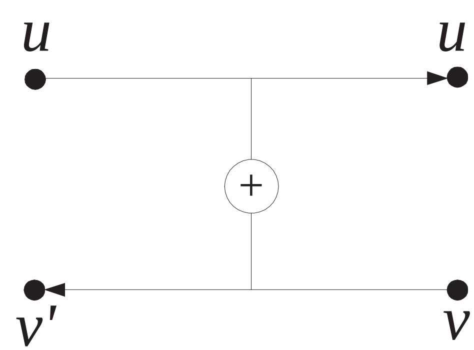

- *At most one* of the 2 schematic edges will be included in a cycle cover.
- There will be $3m$ XOR gadgets, one for each literal.

The Proof: Properties of the XOR Gadget (continued)

Total weight of $-1 - 2 + 6 - 3 = 0$ for cycle covers not entering or leaving it.

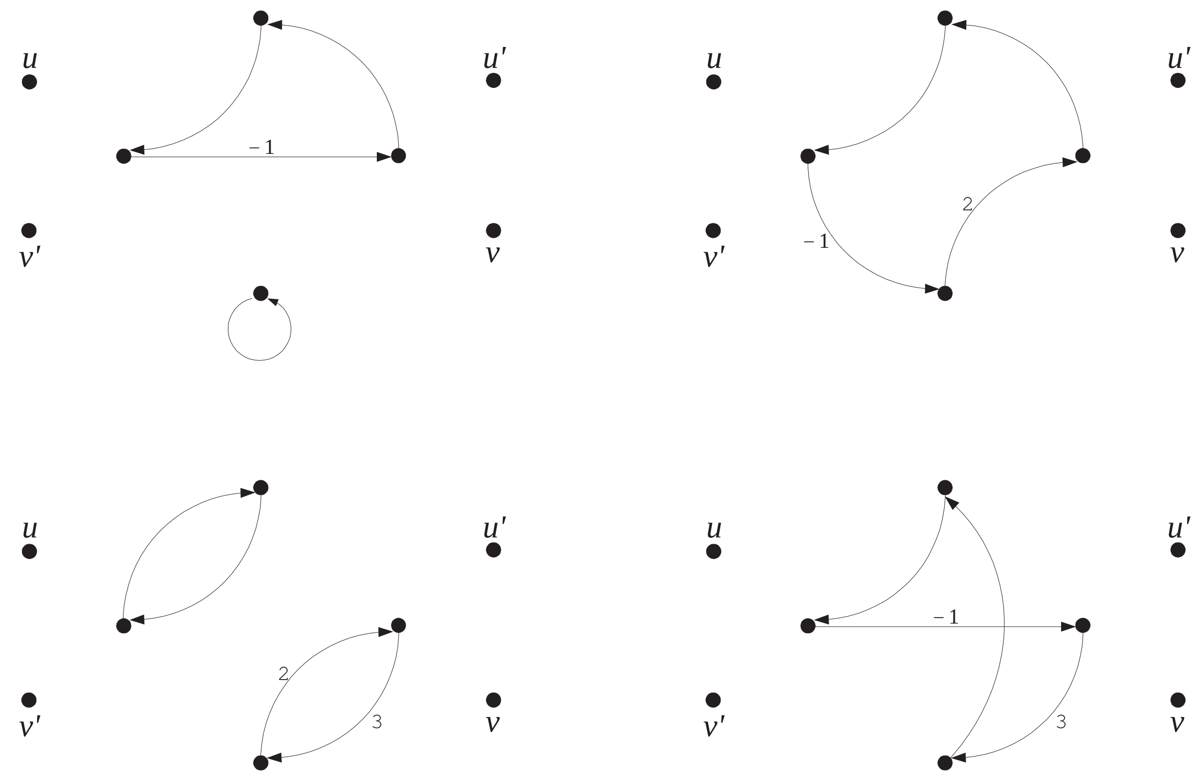

The Proof: Properties of the XOR Gadget (continued)

- Total weight of $-1+1-6+2+3+1=0$ for cycle covers entering at u and leaving at v'.[a]

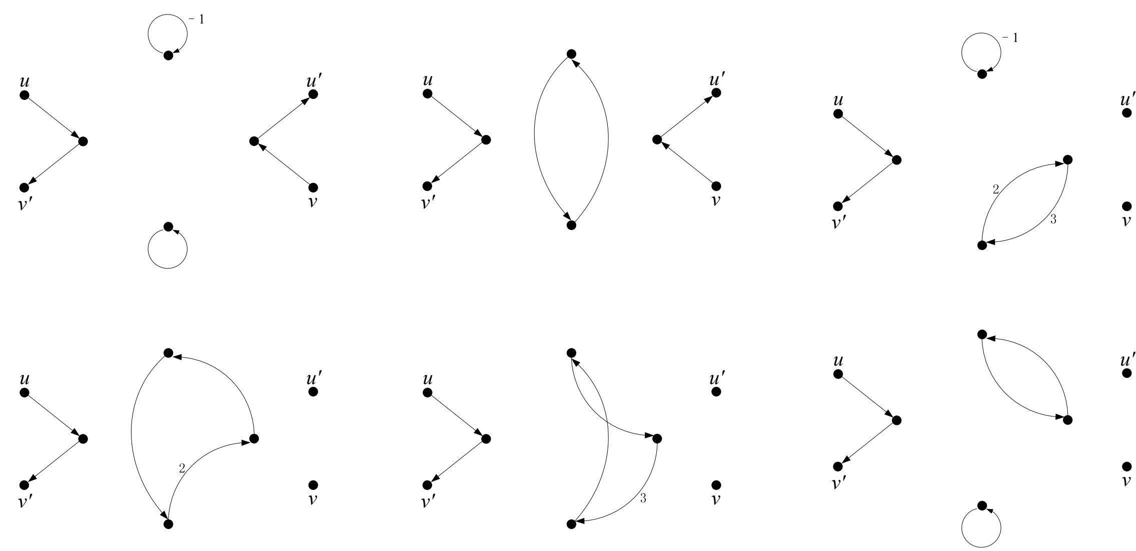

- Same for cycle covers entering at v and leaving at u'.

---

[a]Corrected by Mr. Yu-Tsung Dai (B91201046) and Mr. Che-Wei Chang (R95922093) on December 27, 2006.

The Proof: Properties of the XOR Gadget (continued)

- Total weight of $1 + 2 + 2 - 1 + 1 - 1 = 4$ for cycle covers entering at u and leaving at u'.

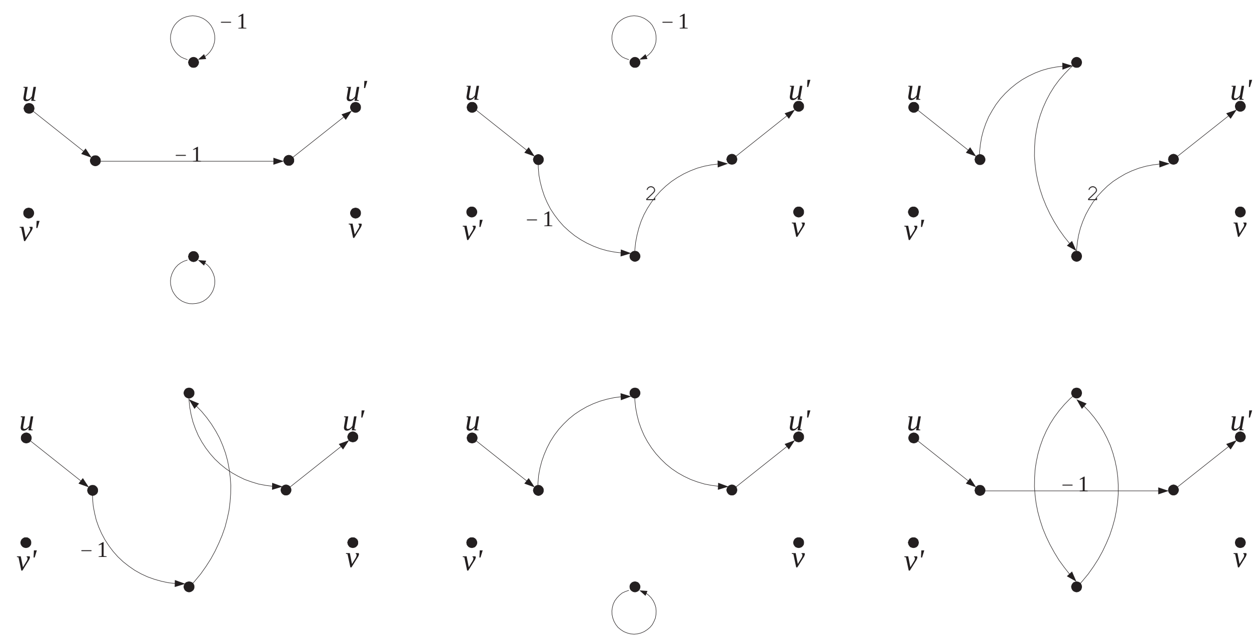

- Same for cycle covers entering at v and leaving at v'.

The Proof: Summary (continued)

- Cycle covers not entering *all* of the XOR gadgets contribute 0 to the cycle count.
  - Let x denote an XOR gadget not entered for a cycle cover c.
  - Now, the said cycle covers' total contribution is

$$
\begin{aligned}
&= \sum_{\text{cycle cover } c \text{ for } H} \text{weight}(c) \\
&= \sum_{\text{cycle cover } c \text{ for } H - x} \text{weight}(c) \sum_{\text{cycle cover } c \text{ for } x} \text{weight}(c) \\
&= \sum_{\text{cycle cover } c \text{ for } H - x} \text{weight}(c) \cdot 0 \\
&= 0.
\end{aligned}
$$

The Proof: Summary (continued)

- Cycle covers entering *any* of the XOR gadgets and leaving illegally contribute 0 to the cycle count.
- For every XOR gadget entered and exited legally, the total weight of a cycle cover is multiplied by 4.
  - With an XOR gadget x entered and exited legally fixed,

$$
\begin{aligned}
& \text{contributions of such cycle covers to the cycle count} \\
& \sum_{\text{cycle cover } c \text{ for } H} \text{weight}(c) \\
= & \sum_{\text{cycle cover } c \text{ for } H - x} \text{weight}(c) \sum_{\text{cycle cover } c \text{ for } x} \text{weight}(c) \\
= & \sum_{\text{cycle cover } c \text{ for } H - x} \text{weight}(c) \cdot 4.
\end{aligned}
$$

The Proof: Summary (continued)

- Hereafter we consider only cycle covers which enter every XOR gadget and leaves it legally.
  - Only these cycle covers contribute nonzero weights to the cycle count.
- They are said to **respect** the XOR gadgets.

The Proof: the Choice Gadget (continued)

- One choice gadget (a schema) for each variable.

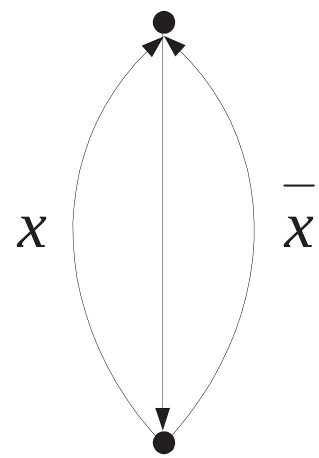

- It gives the truth assignment for the variable.
- Use it with the XOR gadget to enforce consistency.

Schema for $(w \vee x \vee \bar{y}) \wedge (\bar{x} \vee \bar{y} \vee \bar{z})$

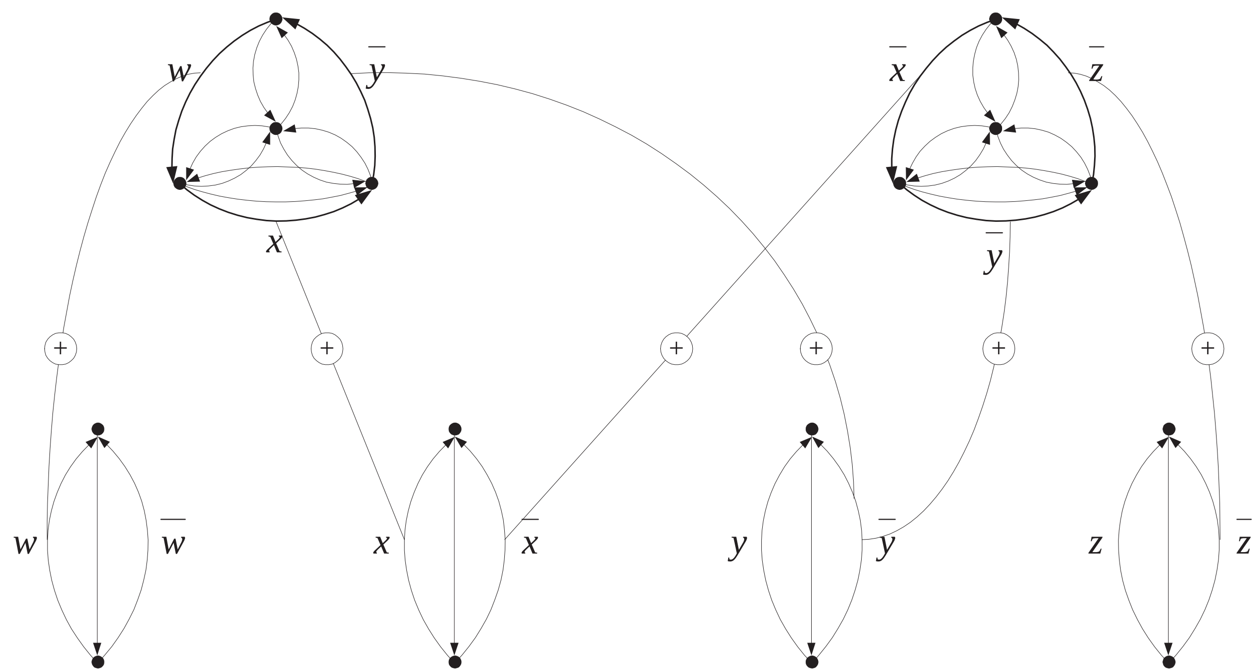

Full Graph $(w \vee x \vee \bar{y}) \wedge (\bar{x} \vee \bar{y} \vee \bar{z})$

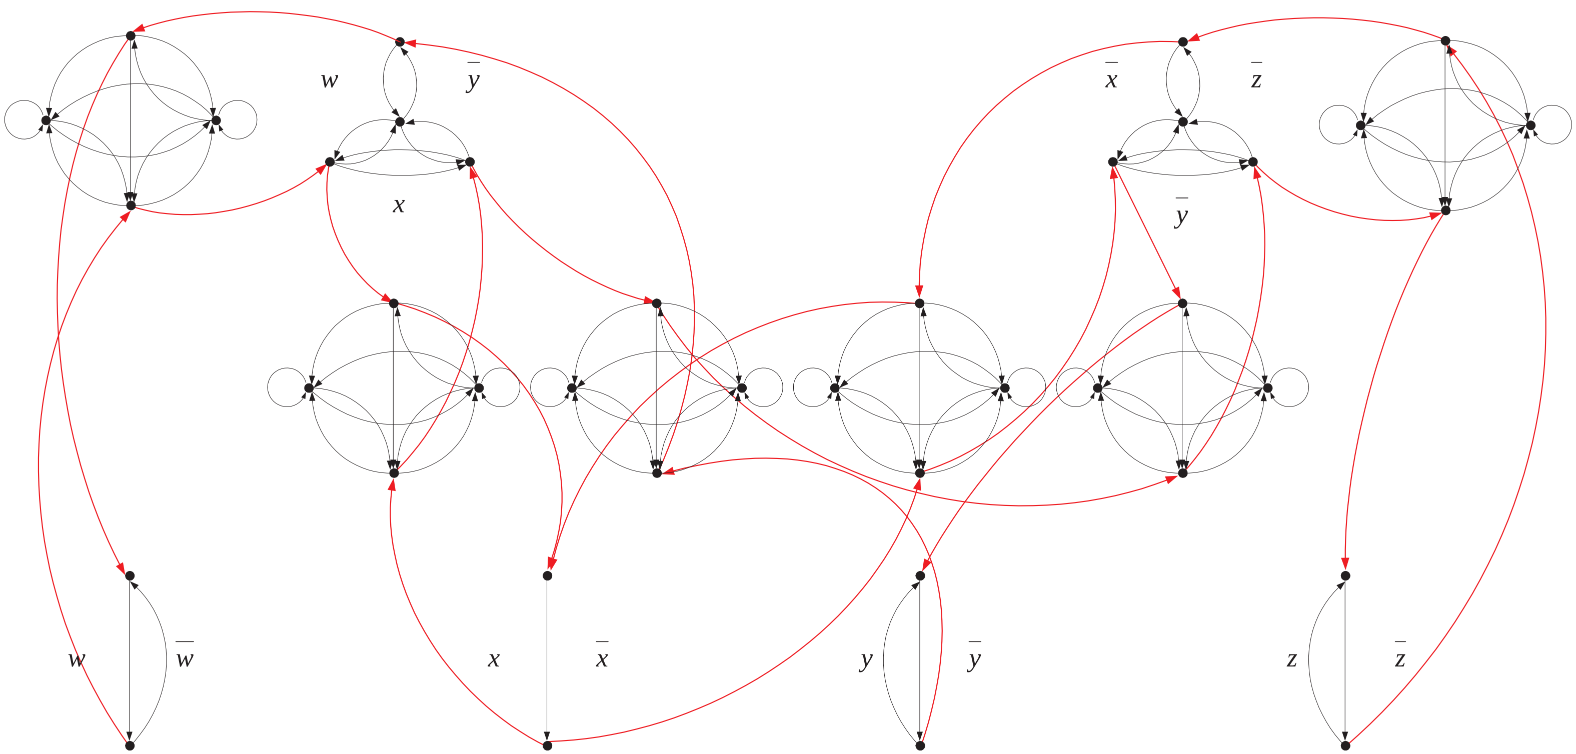

The Proof: a Key Observation (continued)

Each satisfying truth assignment to ϕ corresponds to a schematic cycle cover that respects the XOR gadgets.

$w = 1, x = 0, y = 0, z = 1 \Leftrightarrow$ One Cycle Cover

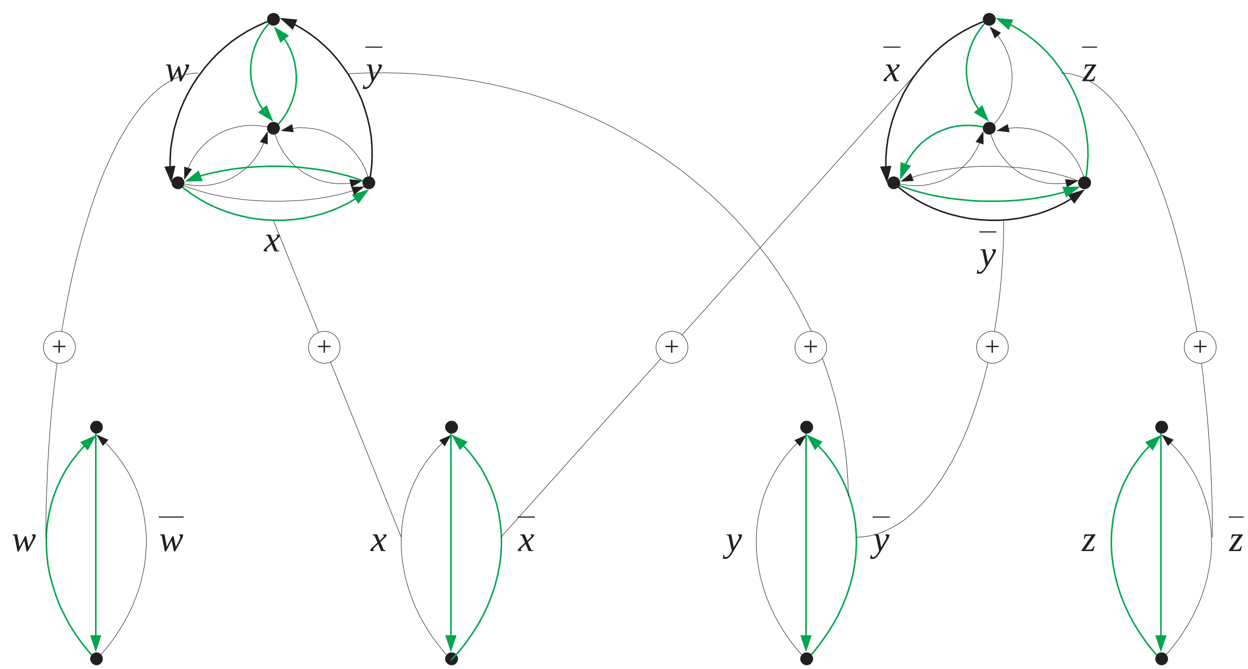

The Proof: a Key Corollary (continued)

- Recall that there are $3m$ XOR gadgets.
- Each satisfying truth assignment to ϕ contributes 4^{3m} to the cycle count $\#H$.
- Hence
$$\#H = 4^{3m} \times \#\phi,$$
as desired.

"$w = 1, x = 0, y = 0, z = 1$" Adds 4^6 to Cycle Count

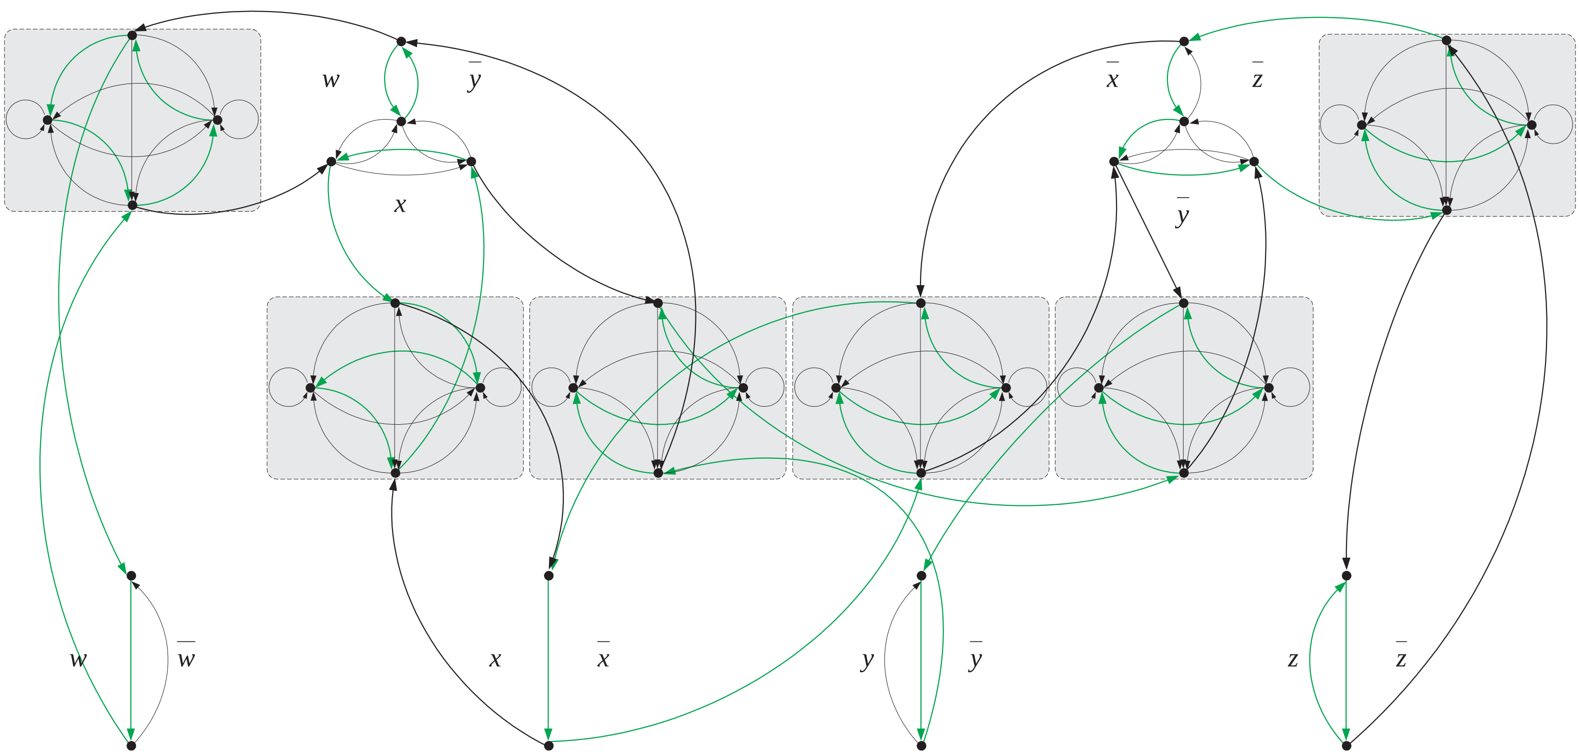

The Proof (continued)

- We are almost done.
- The weighted directed graph H needs to be *efficiently* replaced by some unweighted graph G.
- Furthermore, knowing $\#G$ should enable us to calculate $\#H$ *efficiently*.
  - This done, $\#\phi$ will have been Turing-reducible to $\#G$.[a]
- We proceed to construct this graph G.

[a] By way of $\#H$ of course.

The Proof: Construction of G (continued)

- Replace edges with weights 2 and 3 as follows (note that the graph cannot have parallel edges):

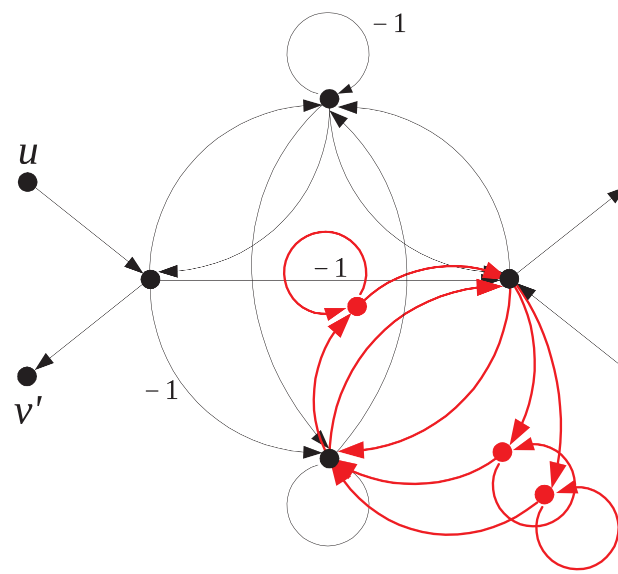

- The cycle count $\#H$ remains *unchanged*.

The Proof: Construction of G (continued)

- We move on to edges with weight -1.
- First, we count the number of nodes, M.
- Each clause gadget contains 4 nodes (p. 761), and there are m of them (one per clause).
- Each revised XOR gadget contains 7 nodes (p. 780), and there are $3m$ of them (one per literal).
- Each choice gadget contains 2 nodes (p. 772), and there are $n \leq 3m$ of them (one per variable).
- So

$$M \leq 4m + 21m + 6m = 31m.$$

The Proof: Construction of G (continued)

- $\# H \leq 2^L$ for some $L = O(m \log m)$.
  - The maximum absolute value of the edge weight is 1.
  - Hence each term in the permanent is at most 1.
  - There are $M! \leq (31m)!$ terms.
  - Hence

$$\begin{aligned} \# H &\leq \sqrt{2\pi(31m)} \left(\frac{31m}{e}\right)^{31m} e^{\frac{1}{12 \times (31m)}} \\ &= 2^{O(m \log m)} \end{aligned} \tag{13}$$

by a refined Stirling's formula.

The Proof: Construction of G (continued)

- Replace each edge with weight -1 with the following:

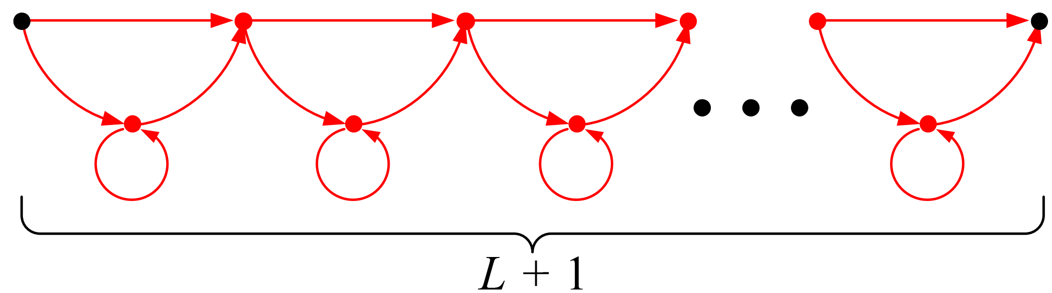

- Each increases the number of cycle covers 2^{L+1}-fold.
- The desired unweighted G has been obtained.

The Proof (continued)

- $\#G$ equals $\#H$ after replacing each appearance -1 in $\#H$ with 2^{L+1}:

$$\begin{aligned} \#H &= \cdots + \overbrace{1 \cdot 1 \cdots (-1) \cdots 1}^{\text{a cycle cover}} + \cdots \,, \\ \#G &= \cdots + \overbrace{1 \cdot 1 \cdots 2^{L+1} \cdots 1}^{\text{a cycle cover}} + \cdots \,. \end{aligned}$$

- Let $\#G = \sum_{i=0}^{n} a_i \times (2^{L+1})^i$, where $0 \leq a_i < 2^{L+1}$.
- Recall that $\#H \leq 2^L$ (p. 782).
- So if we replace -1 by 1, each a_i counts the number of cycle covers with i edges of weight -1 as there is no "overflow."

The Proof (concluded)

- We conclude that

$$\#H = a_0 - a_1 + a_2 - \cdots + (-1)^n a_n,$$

indeed easily computable from $\#G$.

- We know $\#H = 4^{3m} \times \#\phi$ (p. 777).
- So

$$\#\phi = \frac{a_0 - a_1 + a_2 - \cdots + (-1)^n a_n}{4^{3m}}.$$

  – Equivalently,

$$\#\phi = \frac{\#G \bmod (2^{L+1} + 1)}{4^{3m}}.$$

Finis