IN SEARCH OF THE HARD INSTANCES

Albert Atserias

Universitat Politècnica de Catalunya (UPC)
Barcelona

[Cook, 1971]: ACM STOC'71

The Complexity of Theorem-Proving Procedures

Stephen A. Cook

University of Toronto

Summary

It is shown that any recognition problem solved by a polynomial time-bounded nondeterministic Turing machine can be "reduced" to the problem of determining whether a given propositional formula is a tautology. Here "reduced" means, roughly speaking, that the first problem can be solved deterministically in polynomial time provided an oracle is available for solving the second.

...

certain recursive set of strings on this alphabet, and we are interested in the problem of finding a good lower bound on its possible recognition times. We provide no such lower bound here, but theorem 1 will give evidence that {tautologies} is a difficult set to recognize, since many apparently difficult problems can be reduced to determining tautologyhood. By reduced we mean, roughly speaking, that if tautologyhood could be decided instantly

[Cook, 1971]: ACM STOC'71

The field of mechanical theorem proving badly needs a basis for comparing and evaluating the dozens of procedures which appear in the literature. Performance of a procedure on examples by computer is a good criterion, but not sufficient (unless the procedure proves useful in some practical way). A theoretical complexity criterion is needed which will bring out fundamental limitations and suggest new goals to pursue.

Proof Complexity Theorist Dream

1 2 3	2 3 4	-1 -2 -3	-2 -3 -4
1 2 4	2 3 5	-1 -2 -4	-2 -3 -5
1 2 5	2 3 6	-1 -2 -5	-2 -3 -6
1 2 6	2 4 5	-1 -2 -6	-2 -4 -5
1 3 4	2 4 6	-1 -3 -4	-2 -4 -6
1 3 5	2 5 6	-1 -3 -5	-2 -5 -6
1 3 6	3 4 5	-1 -3 -6	-3 -4 -5
1 4 5	3 4 6	-1 -4 -5	-3 -4 -6
1 4 6	3 5 6	-1 -4 -6	-3 -5 -6
1 5 6	4 5 6	-1 -5 -6	-4 -5 -6

Proof Complexity Theorist Dream

1 2 3	2 3 4	-1 -2 -3	-2 -3 -4
1 2 4	2 3 5	-1 -2 -4	-2 -3 -5
1 2 5	2 3 6	-1 -2 -5	-2 -3 -6
1 2 6	2 4 5	-1 -2 -6	-2 -4 -5
1 3 4	2 4 6	-1 -3 -4	-2 -4 -6
1 3 5	2 5 6	-1 -3 -5	-2 -5 -6
1 3 6	3 4 5	-1 -3 -6	-3 -4 -5
1 4 5	3 4 6	-1 -4 -5	-3 -4 -6
1 4 6	3 5 6	-1 -4 -6	-3 -5 -6
1 5 6	4 5 6	-1 -5 -6	-4 -5 -6

REQUIRES 7,904 CLAUSES
TO REFUTE IN RESOLUTION

A SELECTION OF “THEMES” IN PROOF COMPLEXITY

Theme 1 : Lower bounds AND upper bounds

Example: Every Resolution refutation of the pigeonhole principle formulas PHP_n must be of exponential size $2^{\Omega(n)}$

[Haken 1986] answering a question in Cook's 1971 paper

provides lower bounds for the black-box query models of TFNP classes

tight: size $2^{O(n)}$ is an upper bound (but cannot be tree-like!)

Theme 2 : Proof search/Automatability

Given an unsatisfiable CNF formula F

1) find a Resolution refutation of F

2) estimate the Resolution proof length of F

both NP-hard to solve
even **very** approximately
matching subexp. algs.
[A.-Müller 2019]

Theme 3 : Application to analysis of heuristics

- Average-case complexity (e.g., Erdos-Renyi, R3SAT, ...)
- Approximation algorithms (e.g., gap instances)
- Heuristics analysis (e.g., in SAT solving)
- ...

Resolution Inference Rule

given $C \vee x$ and $D \vee \neg x$ infer $C \vee D$

left premise

right premise

resolvent

Tree/Dag Proofs, Size, and Width

empty clause

$$\overbrace{C_1, \dots, C_i, \dots, C_m}^{F}: \qquad \overbrace{D_1, \dots, D_j, \dots, D_k, \dots, D_l, \dots, D_s = \emptyset}^{P}$$

Hypothesis

Left

Right

the proof-graph

Dag-size Tree-size Width Space

Algebraic Proofs

Algebraic proofs:

- Indeterminates x_i and x_i' over a ring ($\mathbb{R}$, $\mathbb{Q}$, $\mathbb{Z}$, $\mathbb{Z}_p$, ...).
- Boolean axioms: $x_i^2 - x_i = 0$ and $x_i + x_i' - 1 = 0$
- Clauses $x_i \vee x_j \vee \neg x_k$ are polynomial eq's $x_i' x_j' x_k = 0$.
- Inferences are polynomial identities.

Sums-of-Monomials Proofs (SOM)
Sums-of-Squares Proofs (SOS)

Let F be a CNF with clauses $C_1,\ldots,C_m$ and variables $x_1,\ldots,x_n$.

$$\sum_i A_i P_i + \sum_k c_k R_k = P$$

a sound proof that F entails $P \geq 0$

clauses of F or Boolean axioms

the "lift" polynomials

non-negative coefficients

SOS: square polynomials P^2
SOM: monomials M

degree : max degree of A_iP_i's and R_k's
monomial size : number of monomials in the A_iP_i's and R_k's
bit size : bit complexity of the proof (the c_k's)

The Point of SOM and SOS Proofs : Adds Counting

GOAL: From PHP derive $1 - \sum_i x_{ih} \geq 0$

at most one pigeon sits in hole h

SOS proof:

$$\sum_{i \neq j} (x_{ih}\, x_{jh})(-1) + \sum_i (x_{ih}^2 - x_{ih})(-1) + \left(1 - \sum_i x_{ih}\right)^2 = 1 - \sum_i x_{ih}$$

hole h exclusivity clauses

boolean axioms

a square

Meets a Classic of SDP : Lovász' Theta

$$\vartheta(G) = \vartheta_3(G) := \max \quad \sum_{u,v} < x_u, x_v >$$

$$\text{s.t.}$$

$$< x_u, x_v > = 0 \qquad \text{for } uv \notin E(G)$$

$$\sum_u < x_u, x_u > \leq 1$$

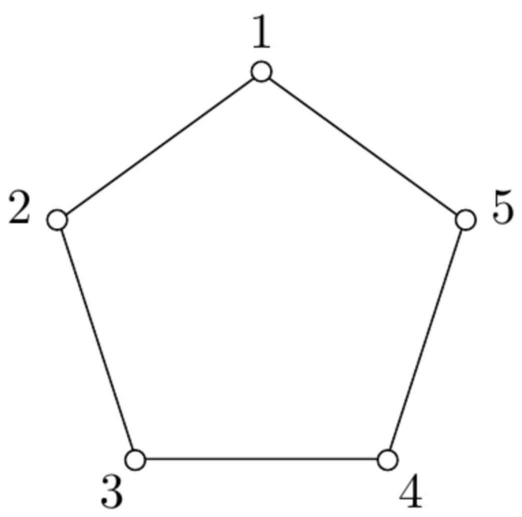

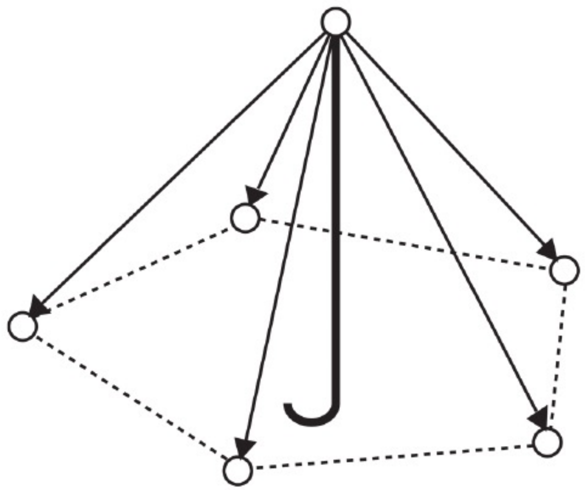

Meets a Classic of SDP : Lovász' Theta

Sandwich Theorem [Lovász 1979]

$$\omega(G) \leq \vartheta(G^c) \leq \chi(G)$$

Theorem [Banks-Kleinberg-Moore 2019]

$\vartheta(G^c) > q$ iff SOS has degree-2 refutation of COL(G,q)

the standard CNF encoding of q-colorability

ANALYSIS OF HEURISTICS

CASE STUDY 1: CLIQUE $K_k \rightarrow G$

CASE STUDY 2: COLORING $G \rightarrow K_q$

CASE STUDY 1:
CLIQUE PROBLEM

The CLIQUE problem

Given a graph G and an integer k
does G have a clique of size k?

Computational complexity of CLIQUE

- NP-complete [Karp'72]
- appears hard on average for $G = G(n, p = n^{-2/(k-1)})$ [Karp'76]
- approximating largest k is NP-hard [Arora-Safra'92, ... PCP ...]
- W[1]-complete when parameterized by k [Downey-Fellows'95]
- requires time $n^{\Omega(k)}$ assuming ETH [Impagliazzo-Paturi'01]
- circuit compexity [Razborov'86, Raz-Wigderson'92, Rossman'10]
- planted clique model [Feige-Krauthgamer'03] [Barak et al.'16]
- etc ...

A common heuristic in practical CLIQUE solvers

1) Greedily properly color the vertices with *few* colors
2) Branch on different color classes
3) Backtrack if “current clique size + remaining colors < k”

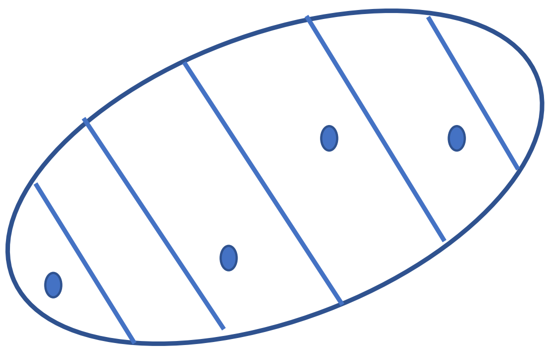

k = 4

More complex heuristics certainly possible (Lovasz theta, etc)

The CLIQUE(G, k) formula

Variables:

$x(i,u)$: "u is the i-th vertex of the clique"

Clauses:

$x(i,1) \vee \ldots \vee x(i,n)$ for i in $[k]$

$\neg x(i,u) \vee \neg x(j,v)$ for i,j in $[k]$ and (u,v) in $V^2 - E$

$G = (V, E)$

$V = [n] = \{1,\ldots,n\}$

k = smaller

Resolution proof complexity of CLIQUE

The trivial upper bound:
The Resolution complexity of CLIQUE(G, k)
is at most $n^{O(k)}$, even for Tree-like Resolution.

Exhaustive enumeration of k-subsets

Question: [Beyersdorff-Galesi-Lauria 2013]
Can one prove that the (general) Resolution complexity of CLIQUE(G, k) can be $n^{\Omega(k)}$?

Motivation 1: Resolution can simulate state of the art practical algorithms

Motivation 2: Answering this seems to require new methods

Answered for tree-like Resolution

Theorem: [BGL 2013]

For $k = O(1)$, the Tree-like Resolution complexity of CLIQUE(G,k) can be $n^{\Omega(k)}$.

Moreover: it is so for $G = G(n, p = n^{-2.01/(k-1)})$ a.a.s.

a "weighted" adversary argument

Question:

What from $G(n, p)$ is really needed to produce the hard instances?

Structure of the lower bound proof

Step 1: If G has a certain property (P),
then Tree-like Resolution complexity
of CLIQUE(G, k) is $n^{\Omega(k)}$.

Step 2: If $G = G(n, p = n^{-2.01/(k-1)})$,
then G has property (P) a.a.s.

Property (P): Rich extension property
Every k/c-subset of vertices
has at least $n^{1-3/c}$ common neighbours

The complete $(k-1)$-partite graph $K(n, k-1)$ has this property too!

... but the (k-1)-colorable graphs are not hard instances, not even for Resolution

Observation: [BGL 2013]

If G is (k-1)-colorable, then

the Resolution complexity of CLIQUE(G, k) is $2^{O(k)}\, n^{O(1)}$.

Compare
Lovasz' Theta's
Sandwich Theorem
$\omega(G) \leq \vartheta(G^c) \leq \chi(G)$

k = 5

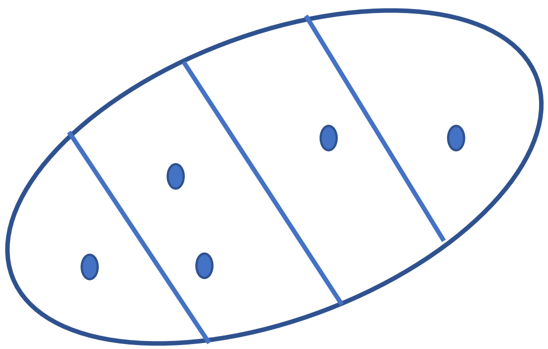

k pigeons
k-1 "meta" holes

Beyond Tree-like Resolution

Theorem: [A.-Bonacina-de Rezende-Lauria-Nordström-Razborov 2019]

For $k = o(n^{1/4})$, the Regular Resolution complexity of $CLIQUE(G,k)$ can be $n^{\Omega(k)}$.

Moreover: it is so for $G = G(n, p = n^{-2.01/(k-1)})$ a.a.s.

Question (again):

What from $G(n, p)$ is really needed to produce the hard instances?

A refined and novel Extension Property (P)

“Clique-Density” Property (P):
Every k/c-set of vertices
has many common neighbours
and
for every set W of vertices for which
every k/cd-set has enough common neighbours in W,
there exists a smallish set S such that
every k/c-set that doesn’t have many common neighbours in W
intersects S at k/cd places.

Sanity check:
Not true in K(n,k-1)!

Lessons learned from CLIQUE

- Resolution complexity brings new perspective into $\omega(G) \leq \mathrm{INT}(G) \leq \chi(G)$.

Could LP-size replace SDP in INT(G) and still get an efficient interpolant?

- A new (convoluted) density property of G(n, p) was identified.

Open: simplify (expander-style?).
Does it hide a new concept?
Can explicit graphs be found?

- **Still open**: Can the (general) Resolution complexity of CLIQUE(G, k) be $n^{\Omega(k)}$? Does Clique-Density suffice?

CASE STUDY 2: COLORING PROBLEM

The COLORING problem

Given a graph G and an integer q
can the vertices of G be q-colored without
monochromatic edges?

Computational complexity of COLORING

- NP-complete even for fixed $q \geq 3$ [Karp'72]
- appears hard on average for $G = G(n, p = 2q \ln(q) / n)$
- approximating $\chi(G)$ is a major problem [... PCP/UGC ...]
- etc ...

The COL(G, q) formula

Variables:

$y(u, i)$: "u is coloured i"

Clauses:

$x(u,1) \vee \ldots \vee x(u,q)$ for u in $[n]$

$\neg x(u,i) \vee \neg x(v,i)$ for (u,v) in E and i in $[q]$

Resolution proof complexity of COLORING

Question: [Beame-Culberson-Mitchell-Moore 2005]

What is the worst-case/average-case

Resolution complexity of COL(G, q) formulas?

Motivation:
Resolution
can simulate
many backtracking
algorithms

Resolution models backtracking algorithms

McDiarmid calculus:

lines: non-q-colorable graphs.

axioms: K_{q+1}

inference rule 1:

if $G \subseteq H$, and G is derived,

then derive H

inference rule 2:

if uv is non-edge of H, and $H + uv$ is derived, and H_{uv} is derived,

then derive H

size of proof is defined as number of inference steps

add edge uv

identify u and v

Resolution models backtracking algorithms

Lemma [BCMM 2005]:

If non-q-colorability of G has Tree-like McDiarmid proof of size S,

then $COL(G, q)$ has Resolution refutation of width $O(q^2 + q \log(S))$.

Theorem [BCMM 2005]:

For fixed $q \geq 3$ and large $G = G(n, p = O(1/n))$,

the Resolution complexity of $COL(G,q)$ is, w.h.p.:

width = $\Omega(n)$

size = $\exp(\Omega(n))$

Prove once for Resolution
apply many times
(to many backtracking algorithms)

COL formulas beyond Resolution

Theorem [Krivilevich-Vu 2002] [Coja-Oghlan 2003]

For fixed $q \geq 3$ and large $G = G(n, p = \Omega(q^2/n))$,

it holds that $\vartheta(G^c) > q$ w.h.p.

Recall
Lovasz' Theta:
$\omega(G) \leq \vartheta(G^c) \leq \chi(G)$
Sandwich Theorem

This gives degree-2 SOS refutations of COL(G, q) at average degree q^2 and beyond

Status : Contrast With Random 3SAT

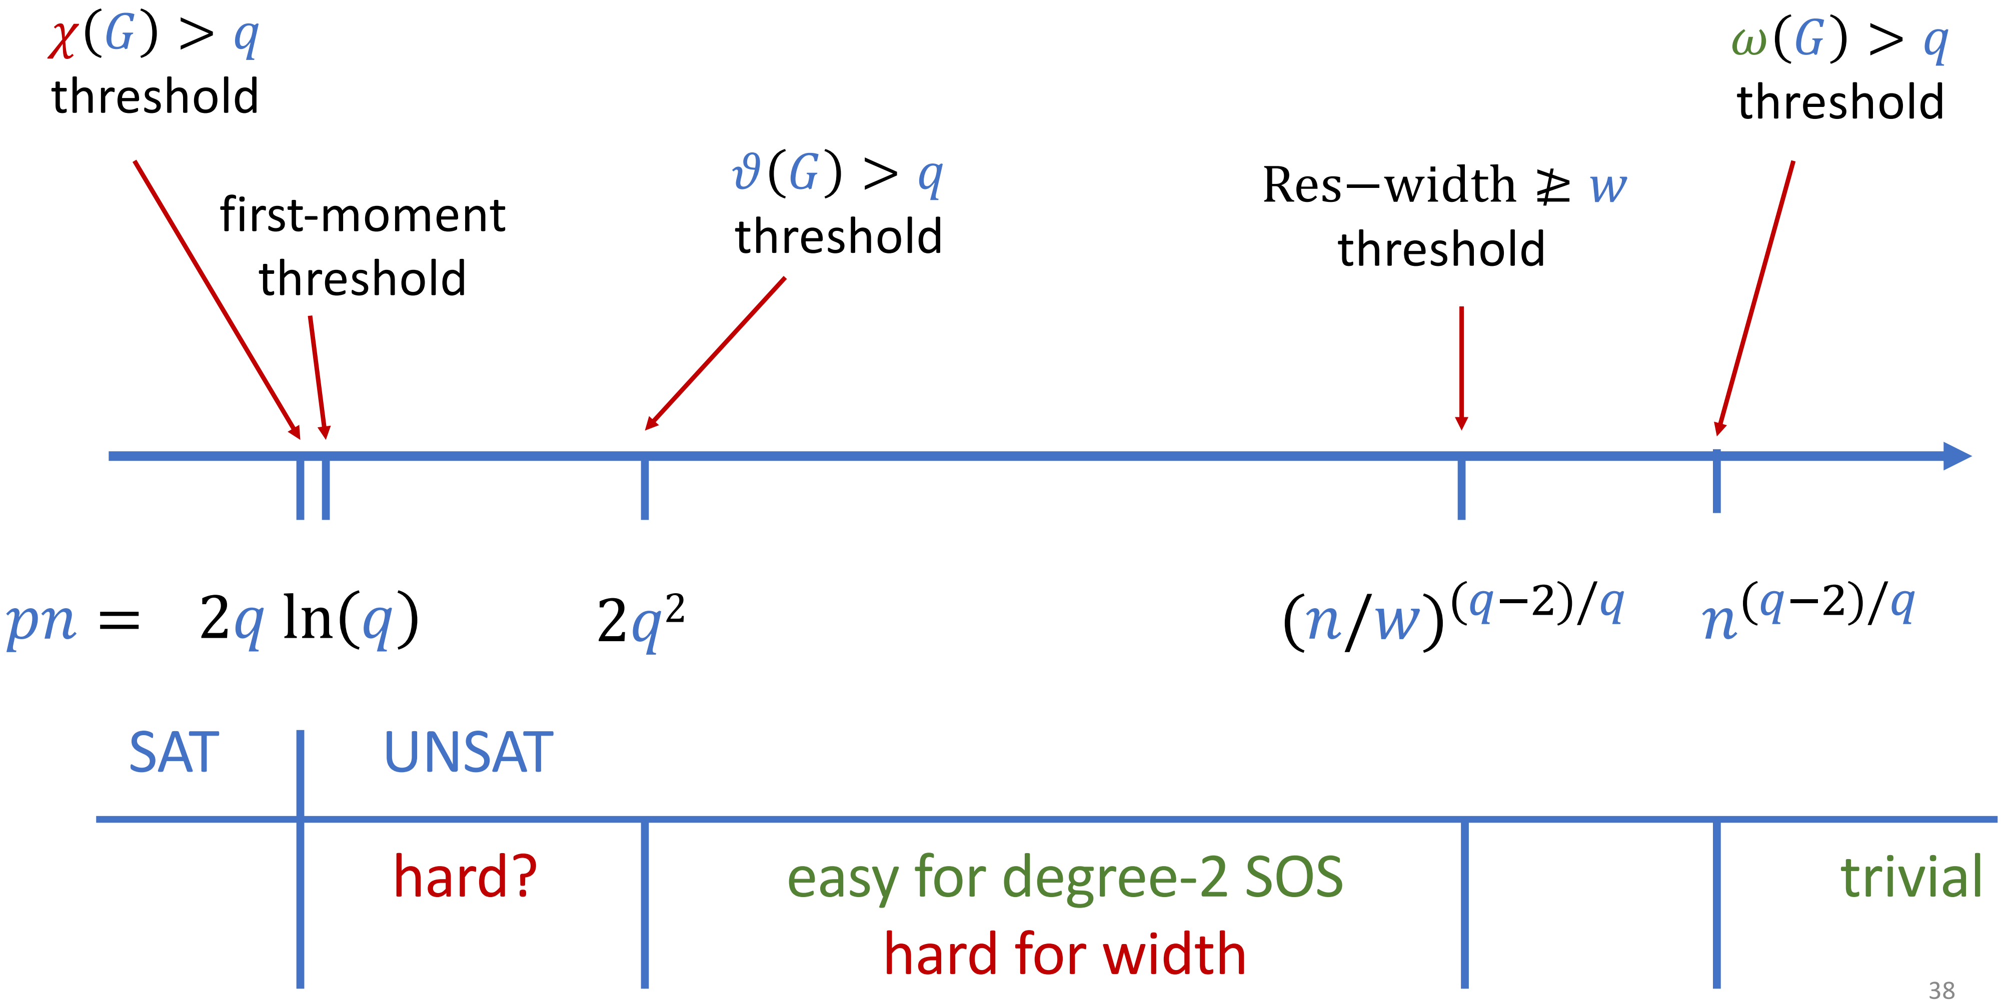

CASE STUDY 2’:
APPROXIMATE GRAPH COLORING

Approximate Chromatic Number

For integers $p \leq q$:

search version

Given a p-colorable graph G,
find a q-coloring of G.

decision version

Given a graph G, output:
YES : if G is p-colorable
NO : if G is not even q-colorable

a promise problem

Computational complexity of approximate χ

3 vs 3 : 3-colorability, NP-complete [Karp'72]

3 vs 4 : NP-complete, PCP Theorem [Khanna-Linial-Safra'00]

3 vs 5 : NP-complete, (PCP +) algebra [Barto-Bulin-Krokhin-Oprsal'21]

3 vs 6 : ?

3 vs q : ?

...

3 vs $n^{1/2}$: in P [Wigderson'83]

NP-complete assuming the d-to-1 Conjecture [Dinur-Mosel-Regev'09]

number of vertices

has been improved many times
current record $0.199 < 1/5$ [Kawarabashi-Thorup'17]

Width-Based Algorithm
Wigderson's algorithm revisited

Much weaker assumption than $\chi(G) \leq 3$ (!)

Fact:

If COL(G, 3) is **not** refutable in width 3,
then G is $O(n^{1/2})$-colorable.

> If COL(G, 3) is not refutable in width 3,
> then G is $O(n^{1/2})$-colorable.

Case 1 : Every u has $d(u) < n^{1/2}$: color greedily as in [W'83].

Case 2 : Some u has $d(u) \geq n^{1/2}$:

enough: as in [W83],
3-color and recurse

Claim: G[N(u) U {u}] is 3-colorable.

Proof:

- If not, then G[N(u)] is not 2-colorable.
- But then COL(G[N(u)], 2) is refutable in width 2: it's a 2-SAT formula.
- So COL(G[N(u) U {u}], 3) is refutable in width 3: add $x_{u,1} \vee x_{u,2} \vee x_{u,3}$
- Hence COL(G, 3) is refutable in width 3. QED

Generalizing further

Thm: [A.-Dalmau'22] Fix ε in $(0,1/2)$.
If COL(G, 3) is not refutable in width $n^{1-2\varepsilon}$,
then G is $O(n^{\varepsilon})$-colorable.

Corollary:
There is an algorithm that solves "3 vs $O(n^{\varepsilon})$" coloring in time $\exp(O(n^{1-2\varepsilon} \log n))$

Beats the naif $\exp(O(n^{1-\varepsilon}))$ bound

If COL(G, 3) is not refutable in width $n^{1-2\varepsilon}$,
then G is $O(n^{\varepsilon})$-colorable.

Case 1 : Some $S \subseteq V$ with $|S| = n^{1-2\varepsilon}$ has $|N(S) \cup S| \geq n^{1-\varepsilon}$.
Case 2 : Every $S \subseteq V$ with $|S| = n^{1-2\varepsilon}$ has $|N(S) \cup S| < n^{1-\varepsilon}$.

Case 1:
a) loop over 3-colorings of G[S],
b) unit propagate to N(S),
c) try to 3-color $G[N(S) \cup S]$,
d) on success:
e) recurse on $G[V-(N(S) \cup S)]$.

Case 2:
a) get n^{ε} such S_i with disjoint $N(S_i) \cup S_i$
b) 3-color each $G[S_i]$, so $G[\cup S_i]$,
c) recurse on $G[V-(\cup S_i)]$.

repeat $\leq$ $\mathbf{n^{\varepsilon}}$ **times**

size: $\mathbf{n - n^{1-\varepsilon}}$

A CHALLENGE

Recall

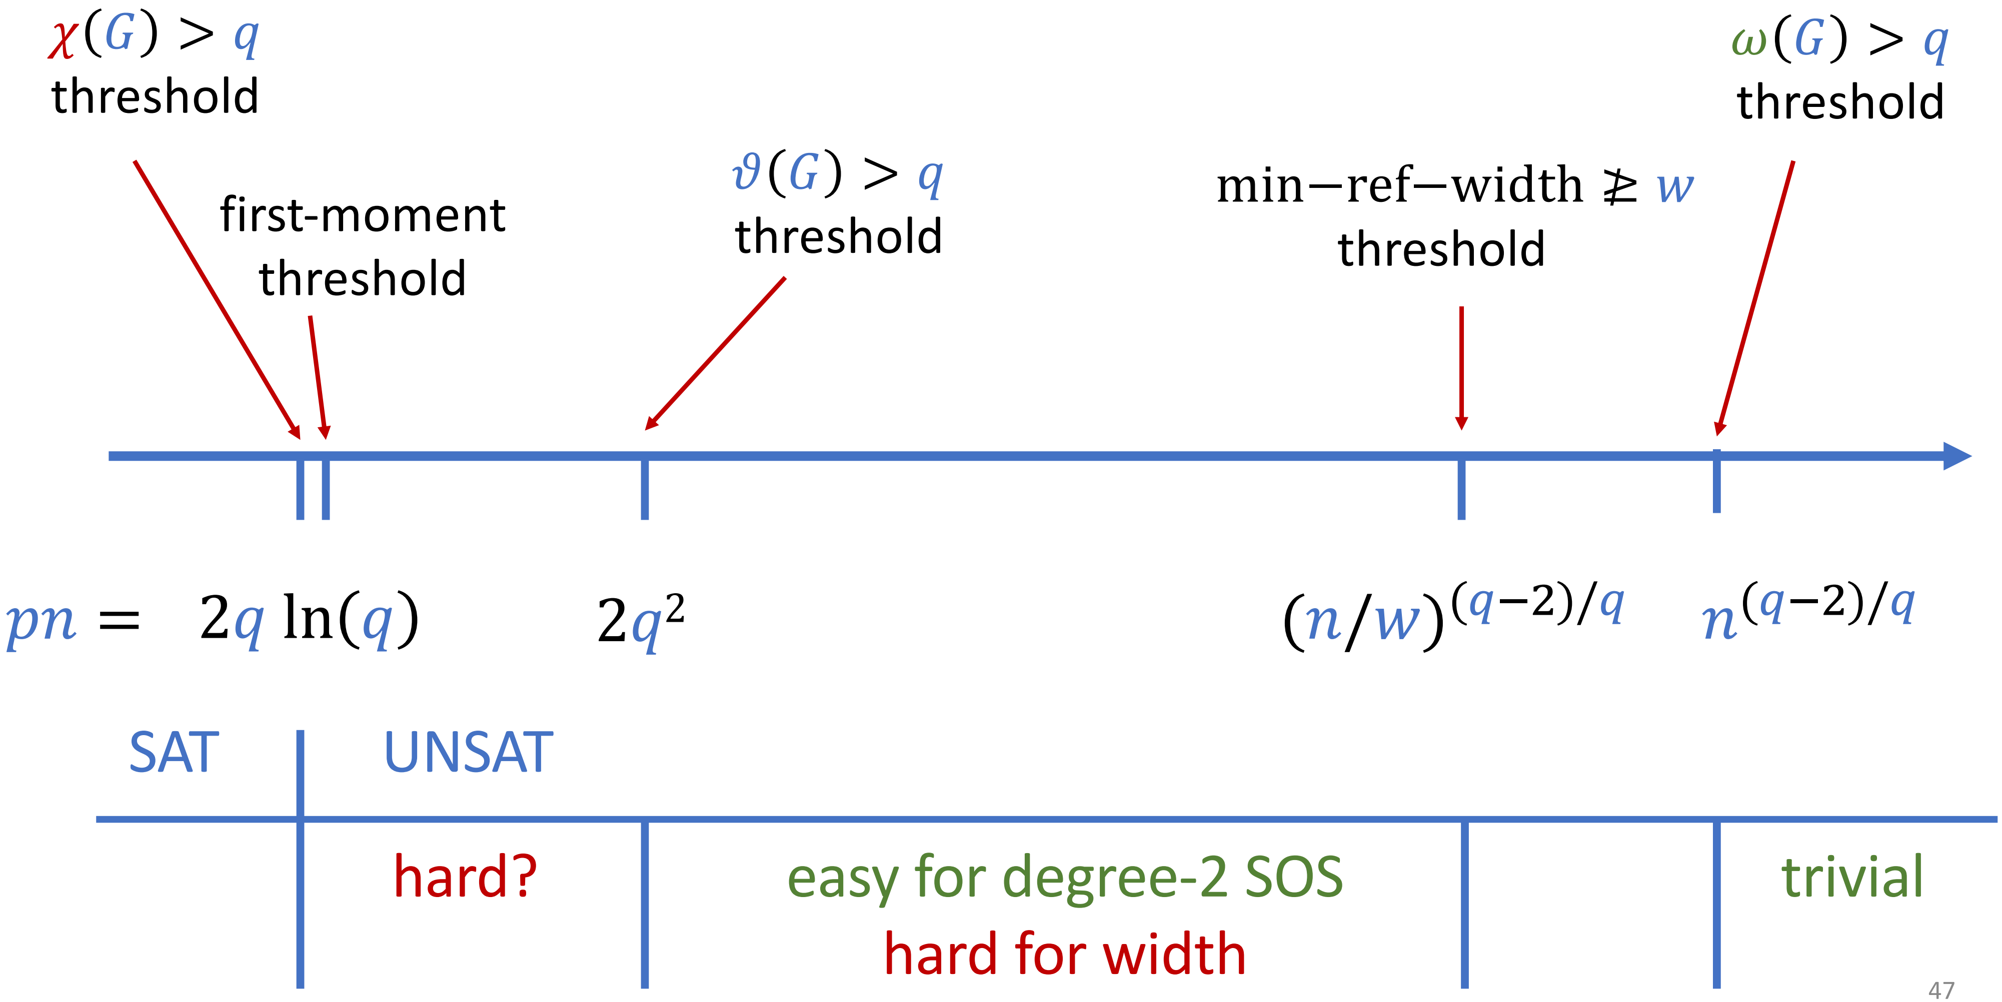

$\chi(G) > q$
threshold
first-moment
threshold
$\vartheta(G) > q$
threshold
$\text{min-ref-width} \ngeq w$
threshold
$\omega(G) > q$
threshold
$pn =$
$2q\ln(q)$
$2q^2$
$(n/w)^{(q-2)/q}$
$n^{(q-2)/q}$
SAT
UNSAT
hard?
easy for degree-2 SOS
hard for width
trivial

Let's Make This Concrete

$q = 3$
$d = 18$
$n =$ large

G(n, d/n)

G_q(n, d/n)

1/2

Left / Right

END