Poly-logarithmic Frege Depth Lower Bounds via an Expander Switching Lemma

Toniann Pitassi (Toronto), Benjamin Rossman (Toronto), Rocco Servedio (Columbia) and **Li-Yang Tan** (TTI-Chicago)

STOC, Boston MA
20 June 2015

Propositional Proof Complexity

Given a universally true statement (a tautology) φ,
what is the length of the shortest proof of φ?

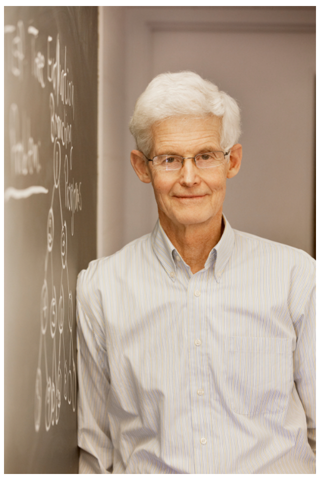

Cook's Program

- NP = coNP iff there is a proof system such that **every short tautology φ has a short proof**.
- We believe NP ≠ coNP, so let's rule this out for increasingly stronger proof systems.

Cook-Reckhow 1979: Let's start with the **Frege proof system**

Remains a flagship open problem of proof complexity today

The Frege Proof System

("Propositional Logic 101 proofs")

Given axioms (Boolean disjunctions), use inference rules to derive trivial tautology

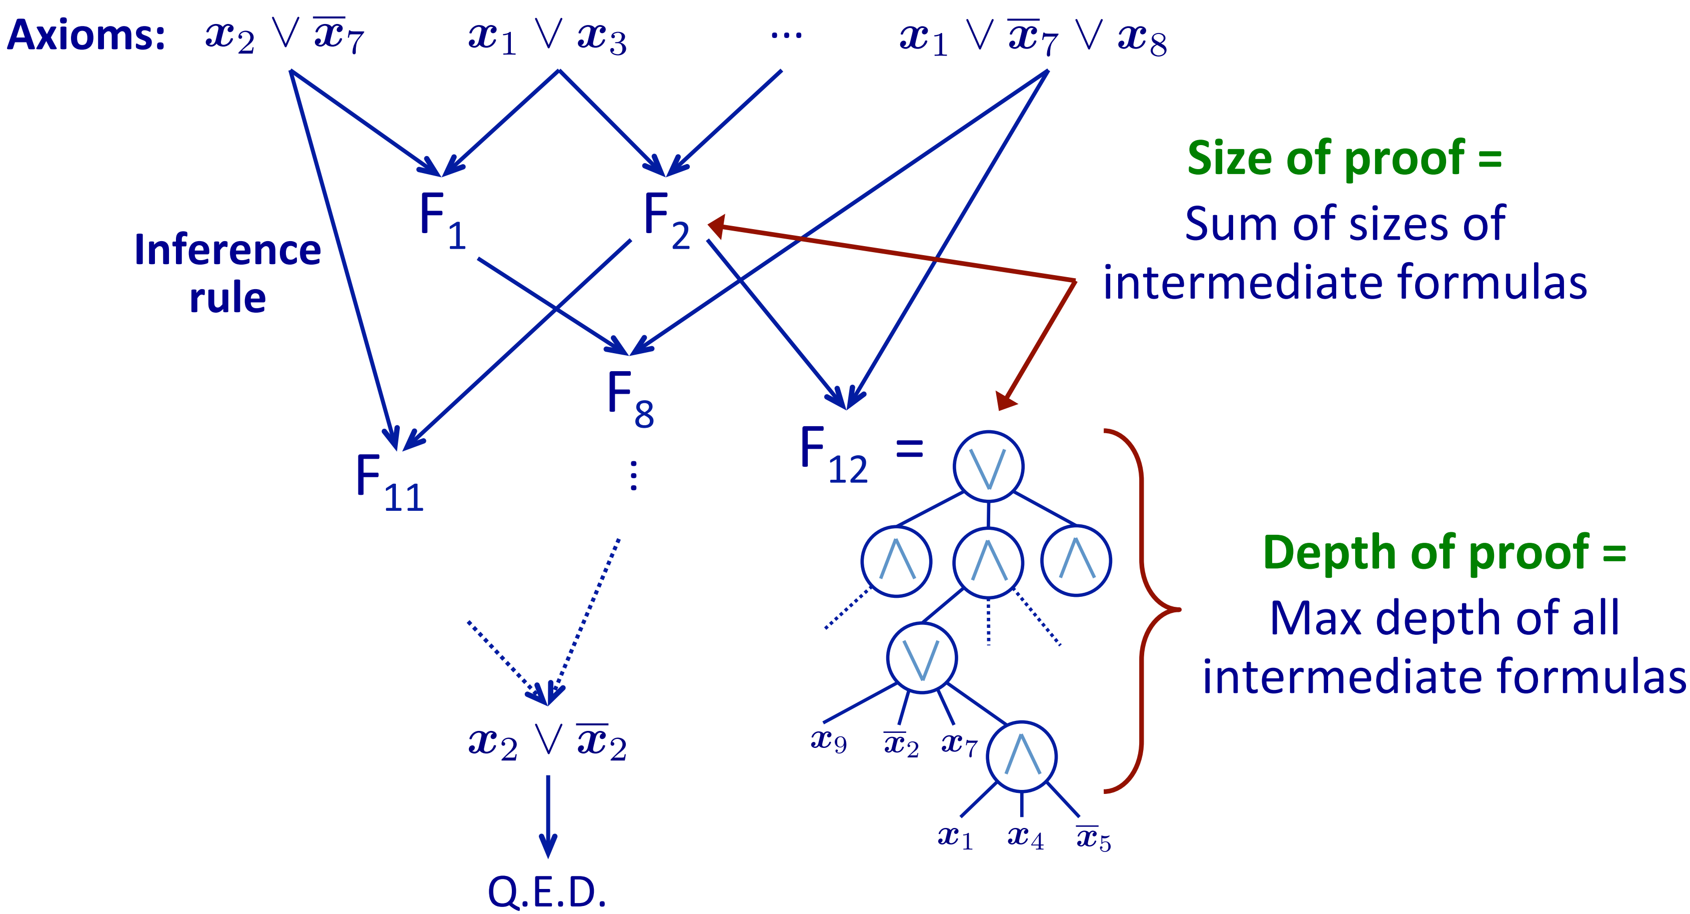

Cook and Reckhow's Challenge (1979):

Super-polynomial size lower bounds against ~~**unrestricted-depth**~~ Frege

Standard fact: suffices to consider **Θ(log n)-depth** Frege

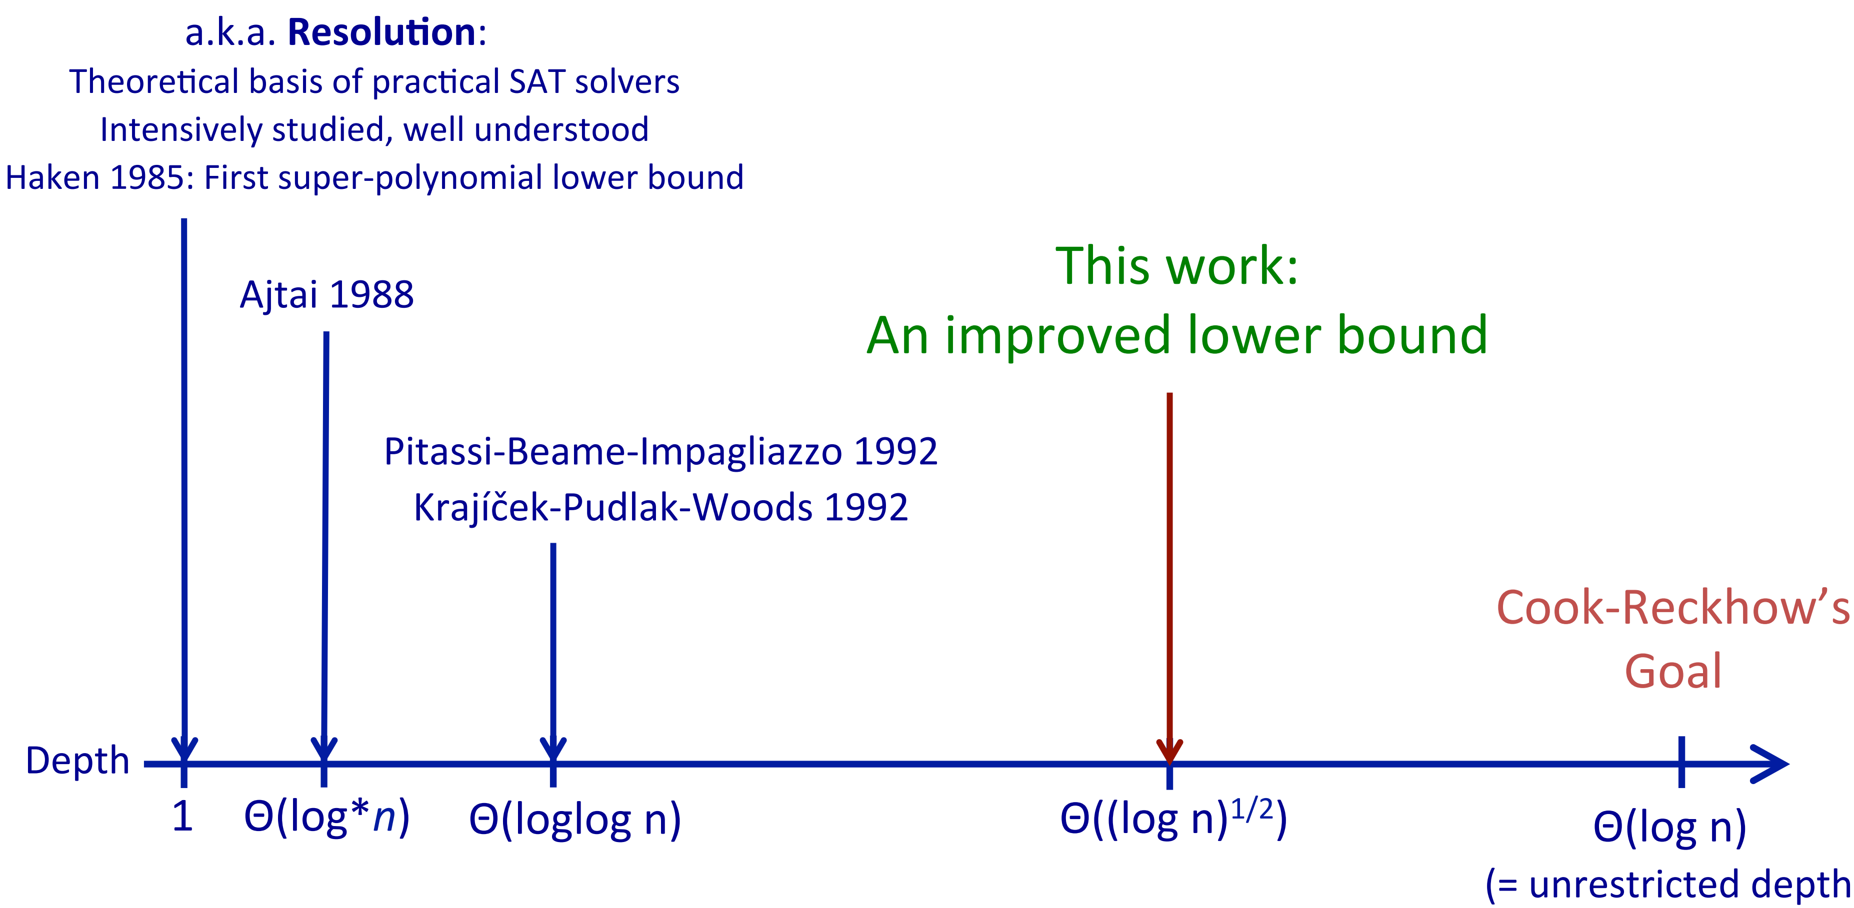

This work (Pitassi-Rossman-Servedio-T 2016)

There is a **linear-size 3CNF tautology φ** such that for all $d = o((\log n)^{1/2})$, every depth-d Frege proof of φ must have **super-polynomial size**.

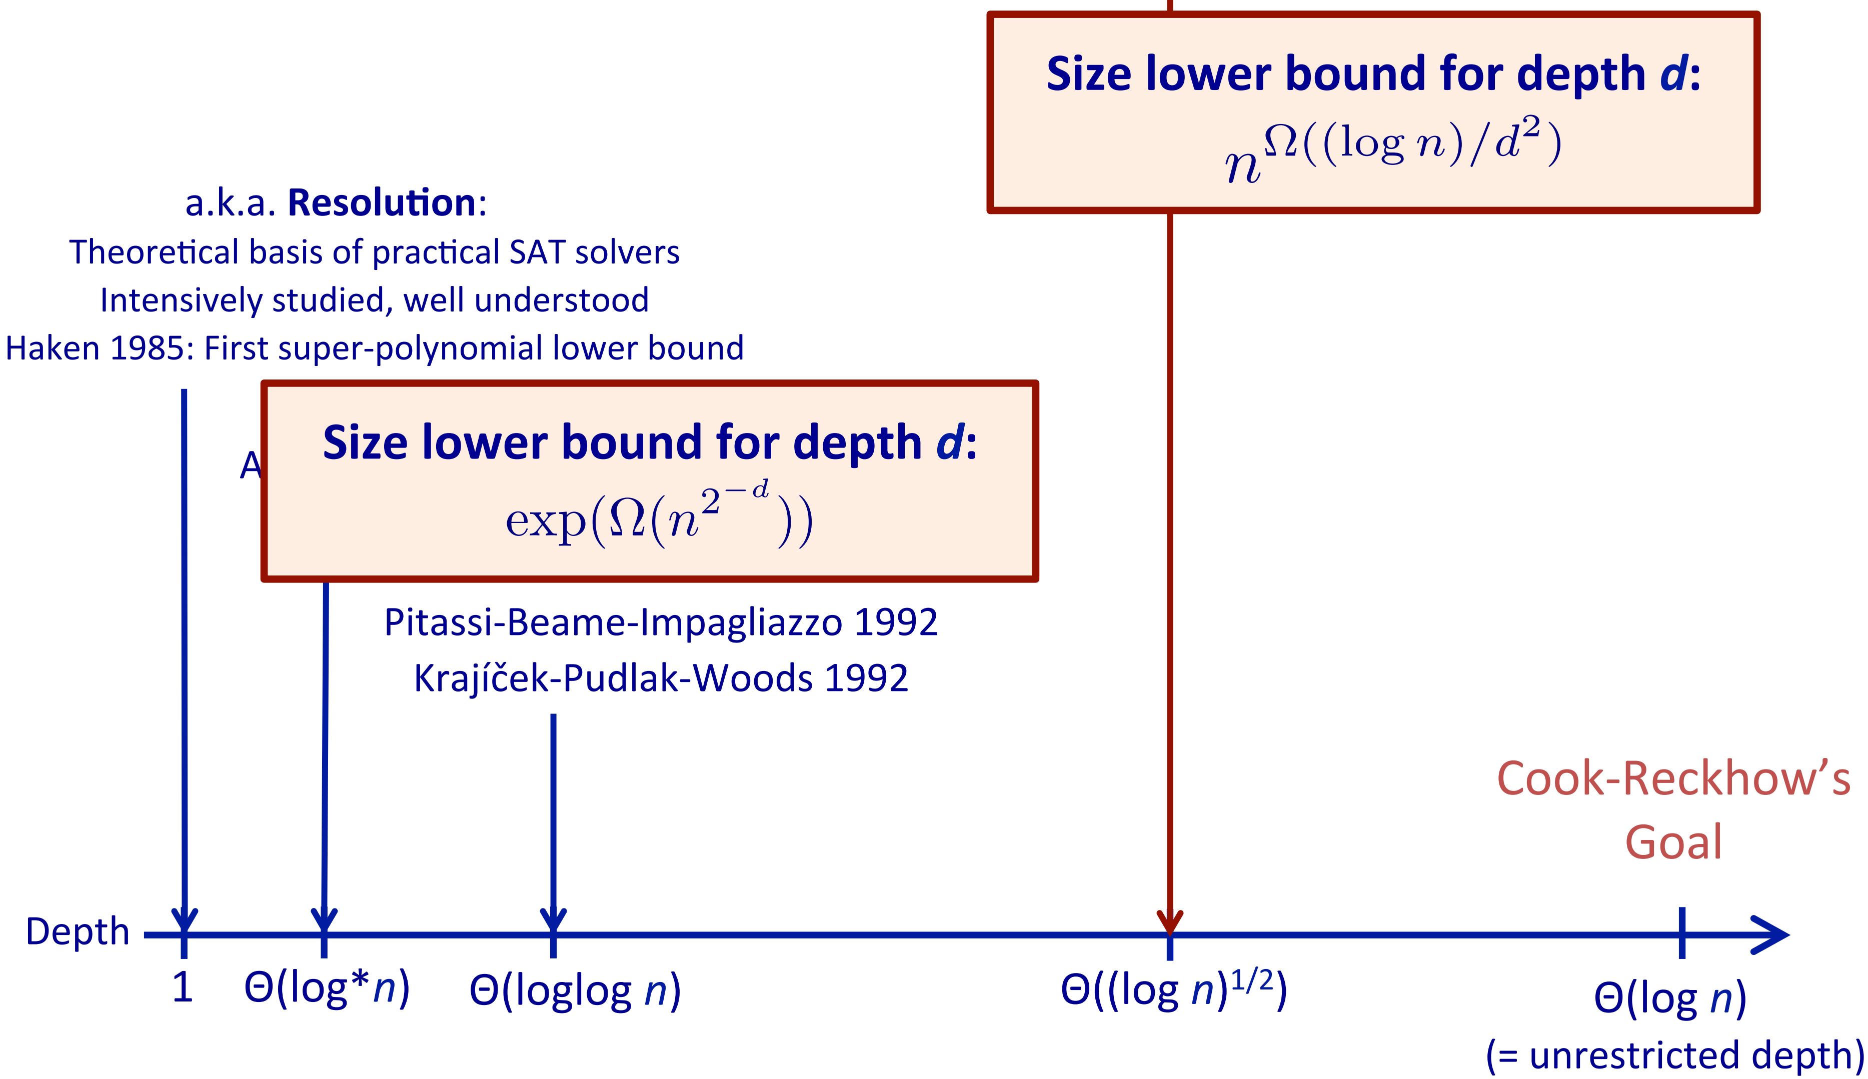

The rest of this talk

- A brief detour into **circuit complexity**
    - PARITY versus AC^0, the role of random restrictions
- Random restrictions in **proof complexity**
- Difficulties faced by previous approaches
- Overcoming these difficulties with random **projections**

PARITY

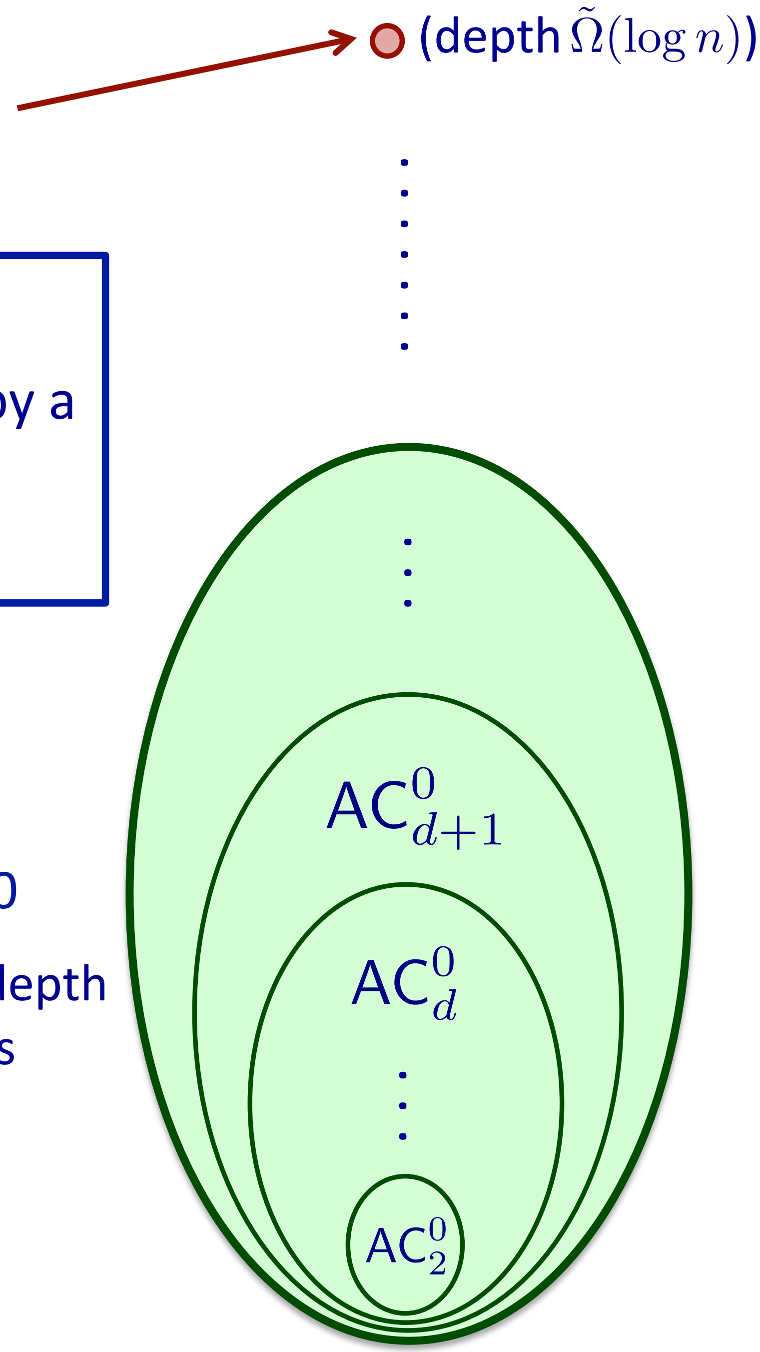

Theorem (FSS81, Ajt83, Yao85, Hås86)

The PARITY function cannot be computed by a constant-depth polynomial-size circuit.

"PARITY not in AC^0"

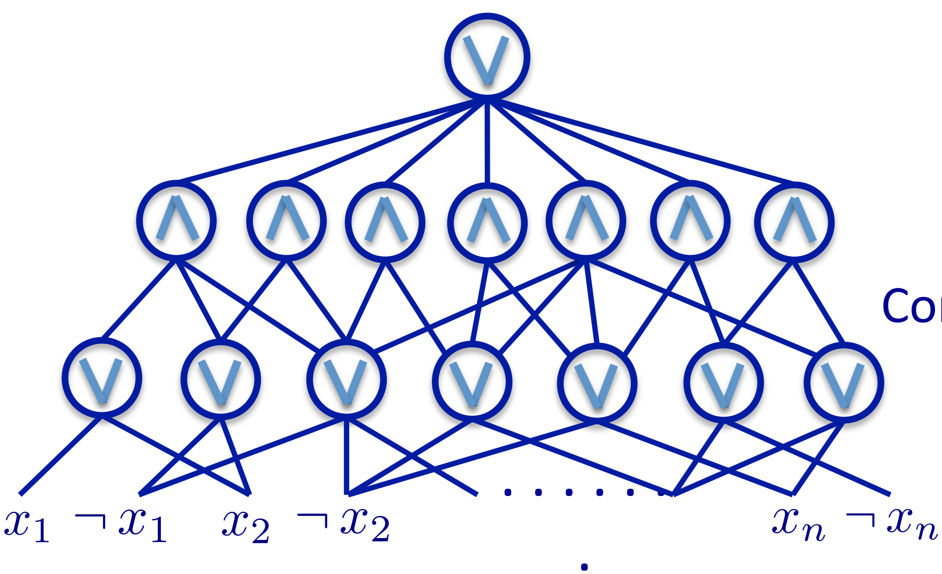

AC^0

Constant-depth circuits

Parity

"remains complex" under random restrictions

Parity on fewer variables

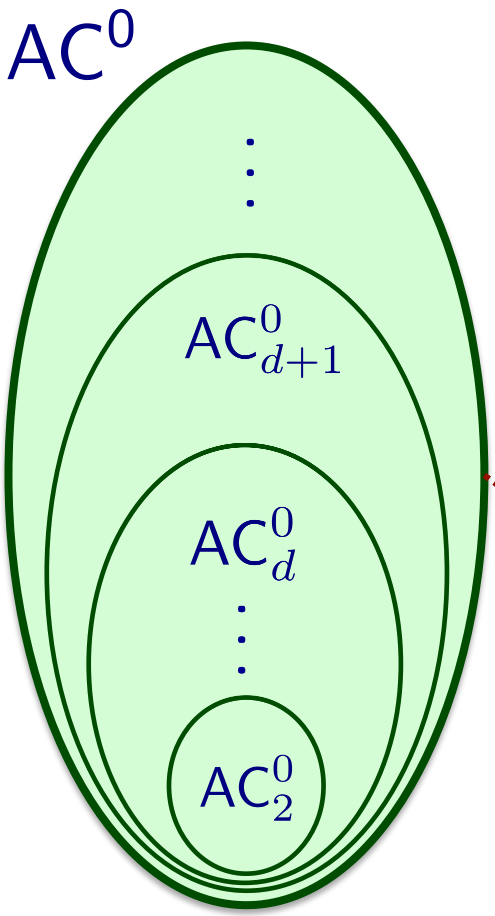

Switching Lemma (FSS81, Ajt83, Yao85, Hås86)

Under random restriction with appropriate *-probability *p*, depth of AC^0 circuits decrease by at least one

collapses to "simple" function
(via Switching Lemma)

DTs

Back to proof complexity

- A brief detour into **circuit complexity**:
  - PARITY versus AC^0, the role of random restrictions
- Random restrictions in **proof complexity**
- Difficulties faced by previous approaches
- Overcoming these difficulties with random projections

Recall our main result:

There is a **linear-size 3CNF tautology φ** such that for all $d = o((\log n)^{1/2})$, every depth-d Frege proof of φ must have **super-polynomial size**.

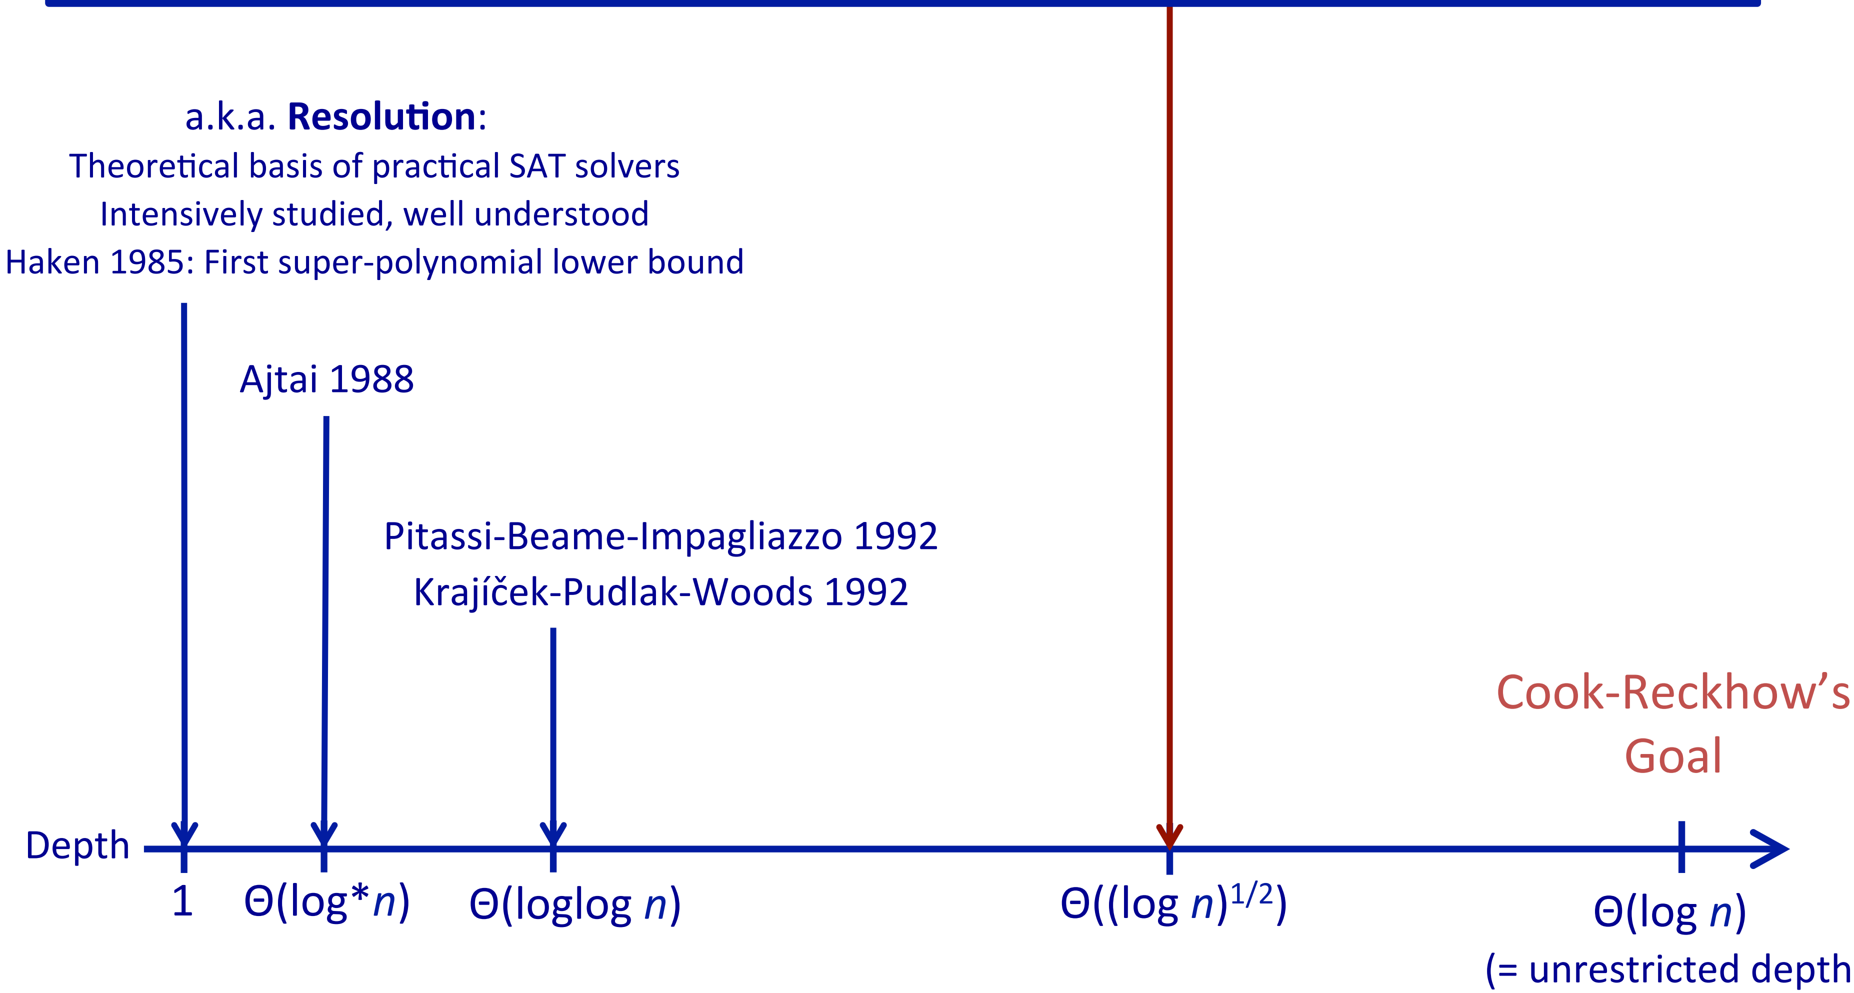

Key difference between our work and previous work

Main technique of previous work:
Random Restrictions

a.k.a. **Resolution**:
Theoretical basis of practical SAT solvers
Intensively studied, well understood
Haken 1985: First super-polynomial lower bound

Our main technique:
Random Projections

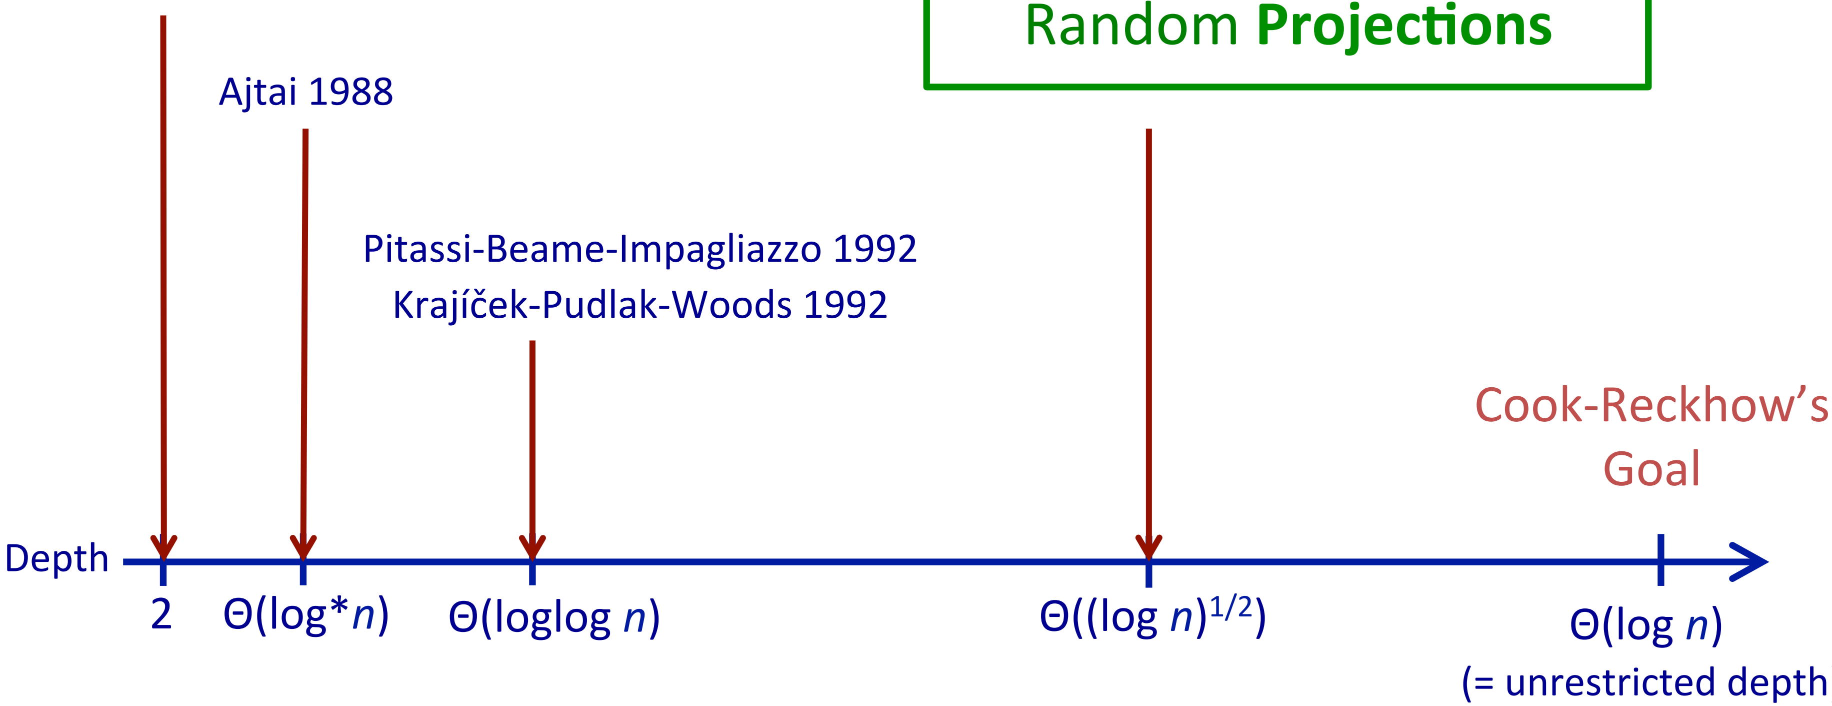

Our hard tautologies φ: Tseitin Tautologies

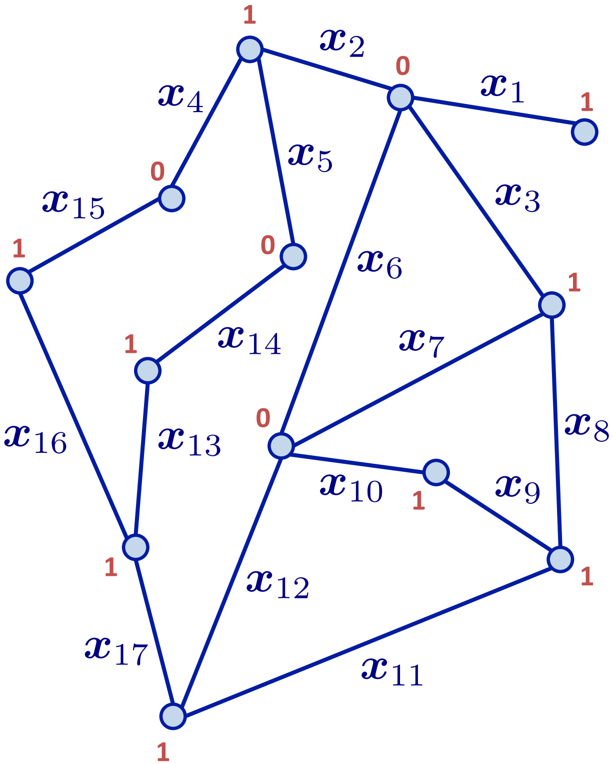

- Underlying graph $G = (V,E)$
- Distinct Boolean variable on each edge
- "Charge" $\alpha : V \to \mathbb{F}_2$ where $\bigoplus_{v \in V} \alpha(v) = 1$

(i.e. sum of charges of all vertices is odd)

Tseitin Tautology ("Handshake lemma")

There is no assignment to edge variables s.t.

$$\bigoplus_{e \sim v} x_e = \alpha(v) \quad \text{for all } v \in V$$

Our hard instances: G = **random 3-regular expander**

Well studied in proof complexity: [..., Urquhart 87, Ben-Sasson 02]

The approach at a high level

(inspired by proof of PARITY not in AC^0)

Tseitin on 3-regular *n*-node expander **vs.** Purported depth-*d* Frege proof

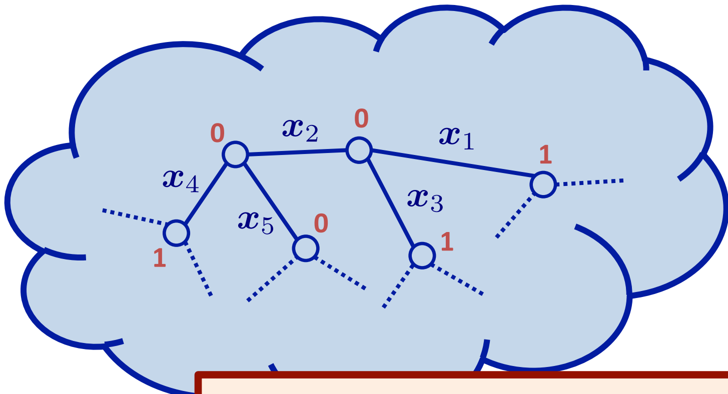

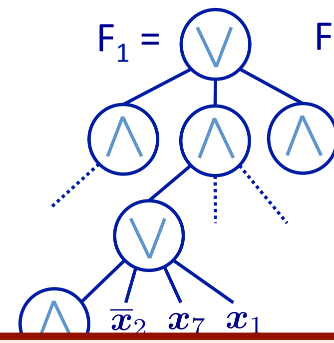

Recall: Every line is a depth-*d* Boolean formula

Main challenge:

Balancing tension between two sides

Random projections give us careful control over this tension

ble” proof

Tseitin on 3-regular *n*’-node expander

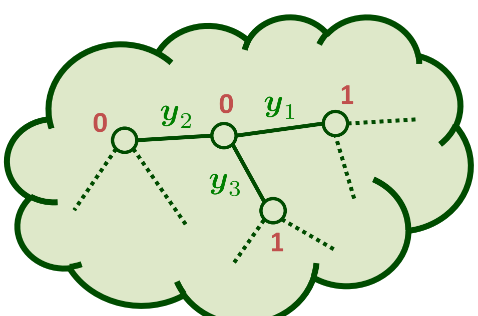

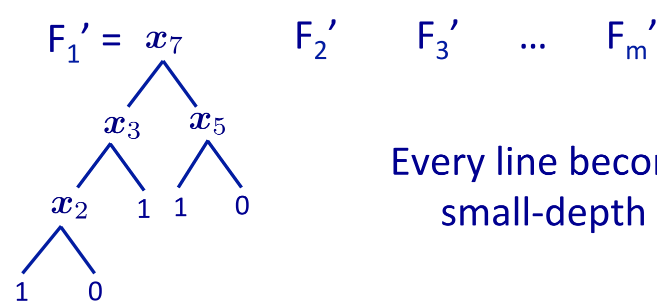

Every line becomes a small-depth DT

Main difficulty of previous approaches:
Keeping hard instances complex under restrictions

Tseitin tautology should not become “too obviously true”

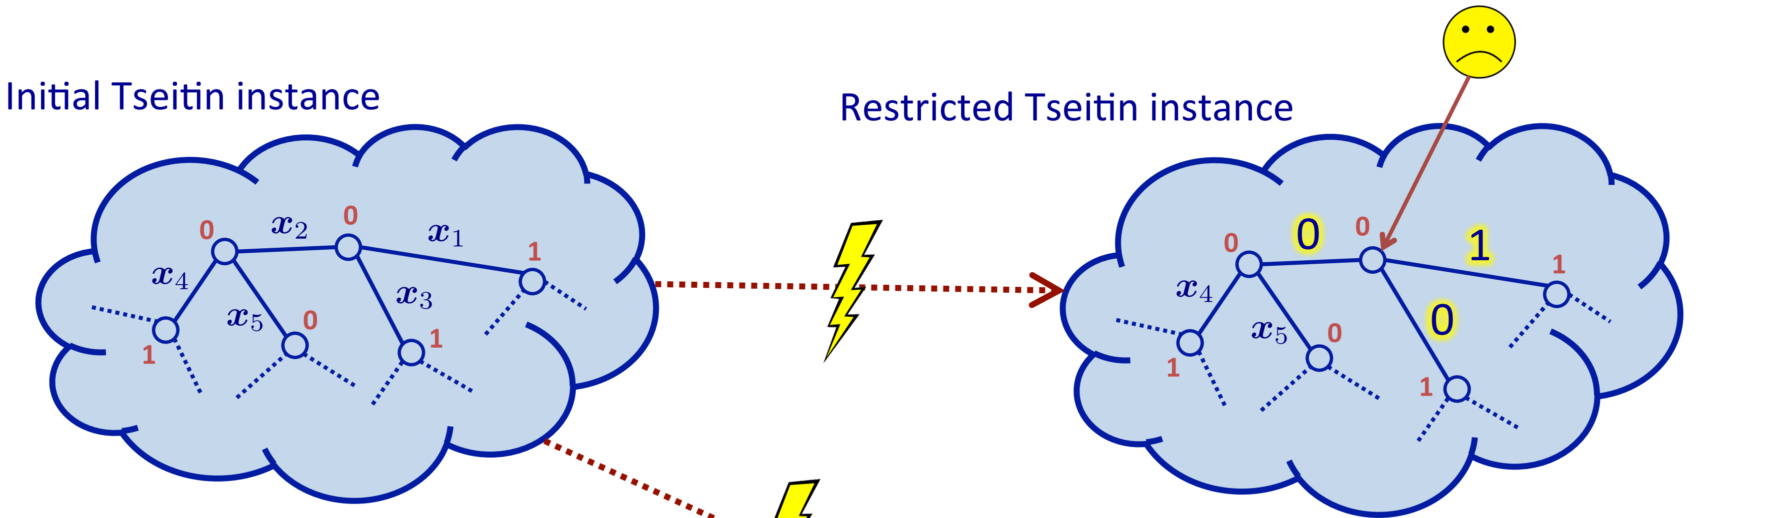

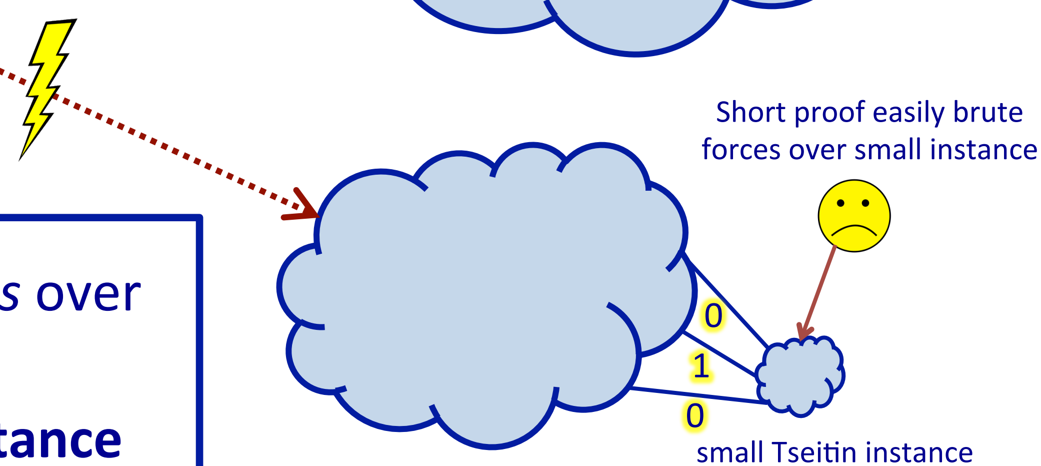

Key advantage of *projections* over restrictions:

Careful control of hard instance

Random Projections

Restriction: Each x_i set to constant or "survives": $x_i \dashrightarrow \begin{cases} 0 \\ 1 \\ x_i \ \text{(denoted *)} \end{cases}$

Projection: Each x_i set to constant or **new formal variable**

$$x_i \dashrightarrow \begin{cases} 0 \\ 1 \\ y_j \end{cases}$$

where $\{ y_1,\ldots,y_m \}$ are new formal variables

Our proof: { y-variables } much smaller than { x-variables }.

Distinct x-variables collide to same y-variable

Our random projection

Step 1: Randomly embed *n*'-node expander in *n*-node expander

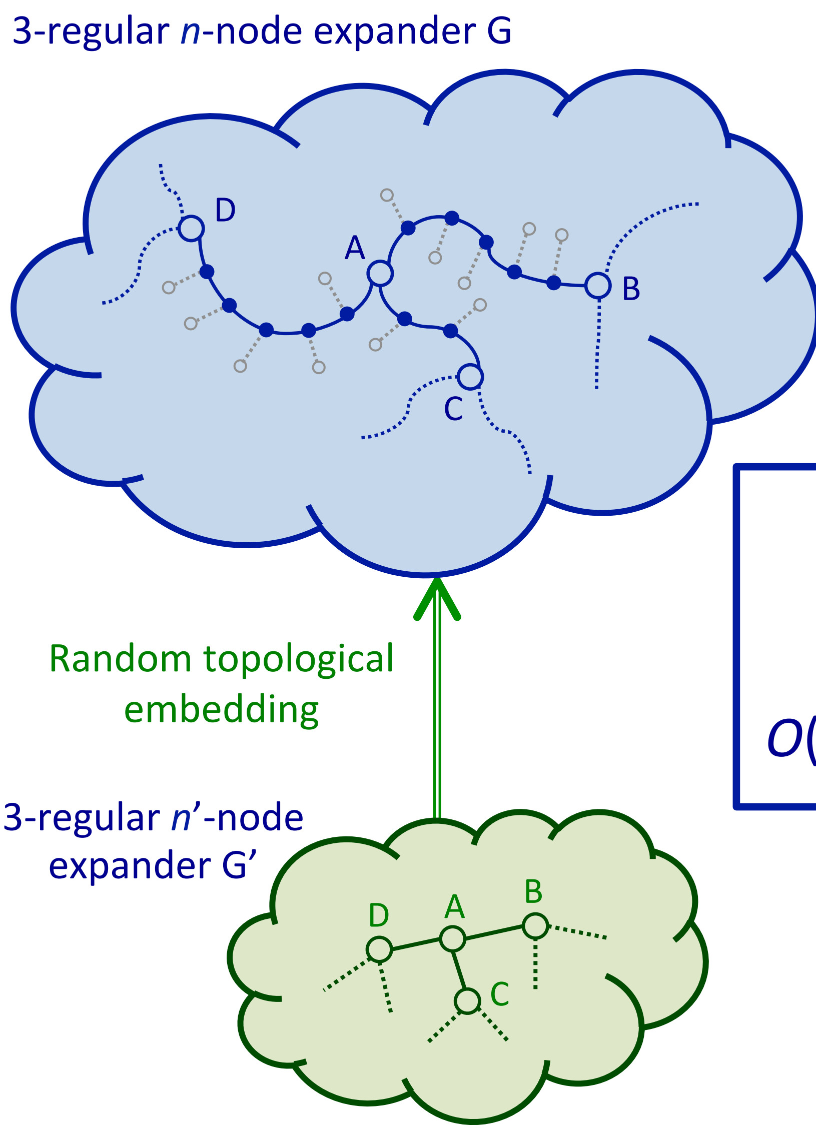

Q: Is this even possible?

What if G does not contain G' as a topological minor?

Theorem (Kleinberg-Rubinfeld 1996)
A bounded-degree *n*-node expander G contains **every graph** G' with $O(n/\text{polylog}(n))$ nodes and edges as a minor.

We build on and extend the algorithmic proof of this theorem.

Our random projection

Step 2: Embedding the Tseitin instance

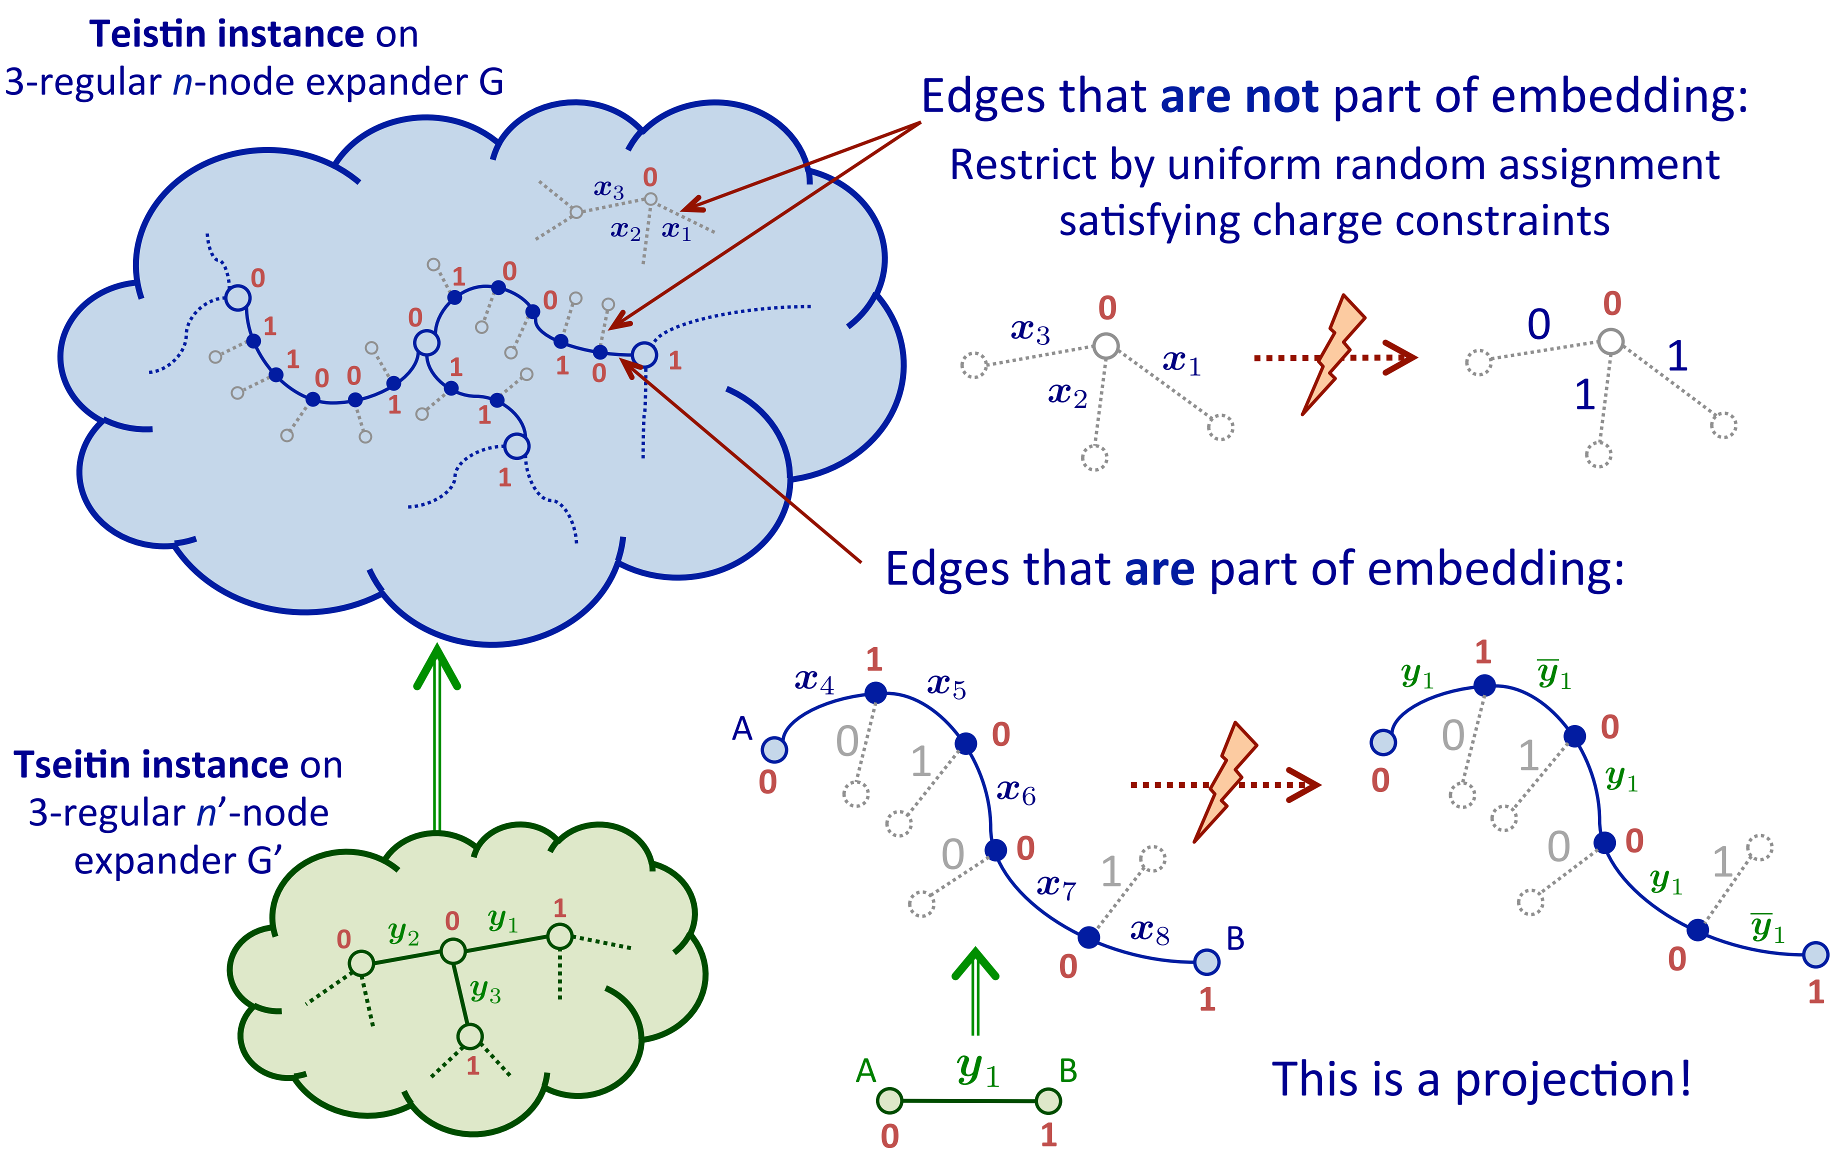

Recap of our approach

Tseitin on 3-regular *n*-node expander **vs.** Purported depth-*d* Frege proof

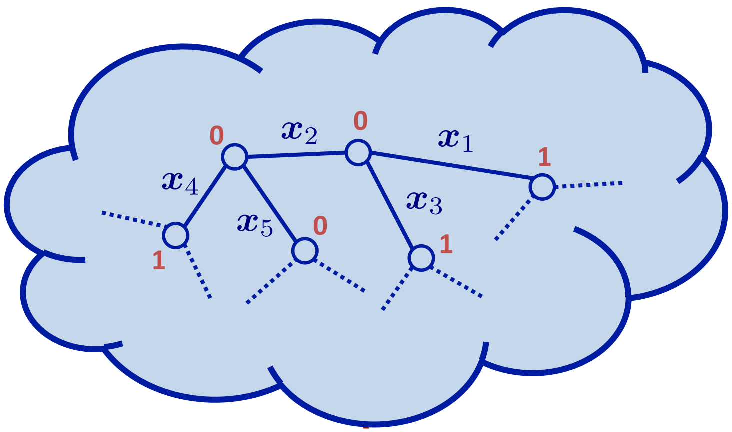

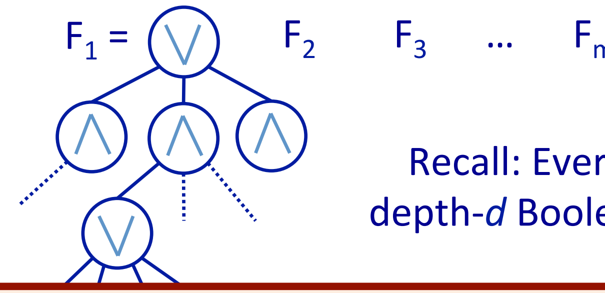

Recall: Every line is a depth-*d* Boolean formula

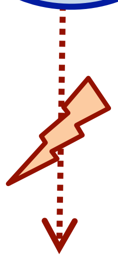

remains "complex"

Main ingredient for this side:

Projection switching lemma for the random projections we just described.

Significant technical challenges

Tseitin on 3-regular *n'*-node expander

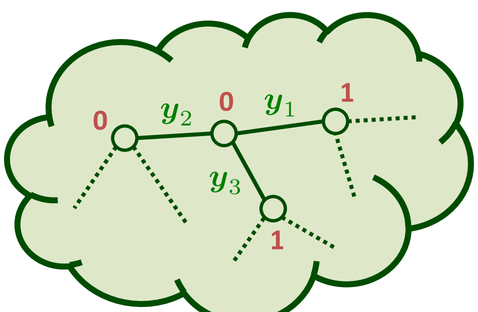

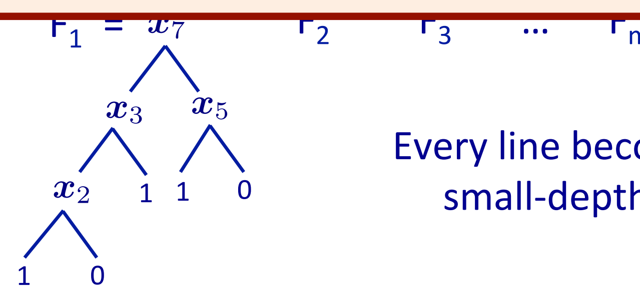

Every line becomes a small-depth DT

Conclusion and open problem

This work (Pitassi-Rossman-Servedio-T 2016)

There is a **linear-size 3CNF tautology φ** such that for all $d = o((\log n)^{1/2})$, every depth-d Frege proof of φ must have **super-polynomial size**.

Size lower bound for depth d:

$$n^{\Omega((\log n)/d^2)}$$

a.k.a. **Resolution**:

Theoretical basis of practical SAT solvers

Intensively studied, well understood

Haken 1985: First super-polynomial lower bound

Conjecture:

$$\exp(\Omega(n^{1/d}))$$

A

Size lower bound for depth d:

$$\exp(\Omega(n^{2^{-d}}))$$

Pitassi-Beame-Impagliazzo 1992

Krajíček-Pudlak-Woods 1992

Cook-Reckhow's Goal

Depth

1   $\Theta(\log^* n)$ $\Theta(\log\log n)$ $\Theta((\log n)^{1/2})$ $\Theta(\log n)$

(= unrestricted depth)

Thank you!