Short XORs for Model Counting: From Theory to Practice

Carla P. Gomes, Joerg Hoffmann,
Ashish Sabharwal, Bart Selman

Cornell University & Univ. of Innsbruck

SAT Conference, 2007
Lisbon, Portugal

Introduction

Problems of interest:

1. **Model counting** (#SAT)
2. Near-uniform **sampling** of solutions

- #P-hard problems, much harder than SAT
- DPLL / local search / conversion to normal forms available but don't scale very well
- Applications to probabilistic reasoning, etc.

A promising new solution approach: *XOR-streamlining*

- **MBound** for model counting
  [Gomes-Sabharwal-Selman **AAAI'06**]
- **XorSample** for sampling [Gomes-Sabharwal-Selman **NIPS'06**]

XOR-Based Counting / Sampling

A relatively simple algorithm:

Step 1. **Add** s uniform random "xor"/parity constraints to F

Step 2. **Solve** with any off-the-shelf SAT solver

Step 3. **Deduce** bounds on the model count of F — **MBound**
or output solution sample of F — **XorSample**

Can boost results further by using exact model counters

Surprisingly good results!

- Counting : Can solve several challenging combinatorial problems previously out of reach
- Sampling : Much more uniform samples, fast

XOR-Based Counting / Sampling

Key Features

- Quick estimates with provable correctness guarantees
- SAT solvers *without modification* for counting/sampling

XOR / parity constraints

- E.g. $a \oplus b \oplus d \oplus g = \text{odd}$ is satisfied iff an odd number of a, b, d, g are True
- Random XOR constraints of **length k** for formula F
  - Choose k variables of F uniformly at random
  - Choose even/odd parity uniformly at random

Focus of this work: ***What effect does length k have?***

Long vs. Short XORs: The Theory

[AAAI'06, NIPS'06]

Full-length XORs (half the number of vars of F)

- Provide provably accurate counts, near-uniform samples

  (both lowerbounds and upperbounds for model counting)
- Limited use: often do not propagate well in SAT solvers

Short XORs (length ~ 8-20)

- Consistently much faster than full-length XORs
- Provide provably correct lowerbounds

  but no upperbounds / good samples
- Quality issue: in principle, can yield very poor lowerbounds

Nevertheless, can short XORs provide good results in practice?

Main Results

Empirical study demonstrating that

- Short XORs often surprisingly good on structured problem instances
  - Evidence based on the fundamental factor determining quality: the *variance* of the random process
- Variance drops drastically in XOR length range ~ 1-5
- Initial results: required XOR length related to (local and global) "backbone" size

What Makes Short XORs Different?

Key difference: **Variance** of the residual model count

Consider formula F with 2^{s^*} solutions. Add s XORs.

- Let X = residual model count
- Same expectation in both cases: $E[X] = 2^{s^*-s}$

- Variance for full XORs : provably low! ($Var[X] \leq E[X]$)
  - Reason: pairwise-independence of full-length XORs
- Variance for short XORs : can be quite high
  - No pairwise-independence

Is variance truly very high in structured formulas in practice?

Experimental Setup

- **Goal**: Evaluate how well short XORs behave for model counting and sampling
  - Comparison with the ideal case: full-length XORs
  - Short XORs clearly favorable w.r.t. time
  - Our comparison w.r.t. ***quality*** of counts and samples
- **Evaluation object**: variance of the residual count
  - Directly determines the quality
  - Compared with the ideal variance (the “ideal curve”) computed analytically
- 1,000 - 50,000 instances per data point

The Quantity Measured

Let X = residual model count after adding s XORs

Must normalize X and s appropriately to compare across formulas F with different #vars and #solns!

Two tricks to make comparison meaningful:

1. Plot variance of **normalized count**: $X' = X / 2^{s^*}$
   - E[X'] = 1 for every F

2. Use $s = s^* - c$ for some *constant* c
   - Var[X'] approaches the same ideal value for every F
   - c : **constant number of remaining XORs**

Experiments: The Ideal Curve

- When variance of X' is plotted with c remaining XORs, can prove analytically

$$\text{ideal-Var} = 2^{-c}$$ for full-length XORs

- We plot sample standard deviation rather than variance

$$\text{ideal-s.s.d.} = \text{sqrt}(2^{-c})$$

For what XOR length does s.s.d.[X'] approach ideal-s.s.d.?

Latin Square Formulas, Order 6

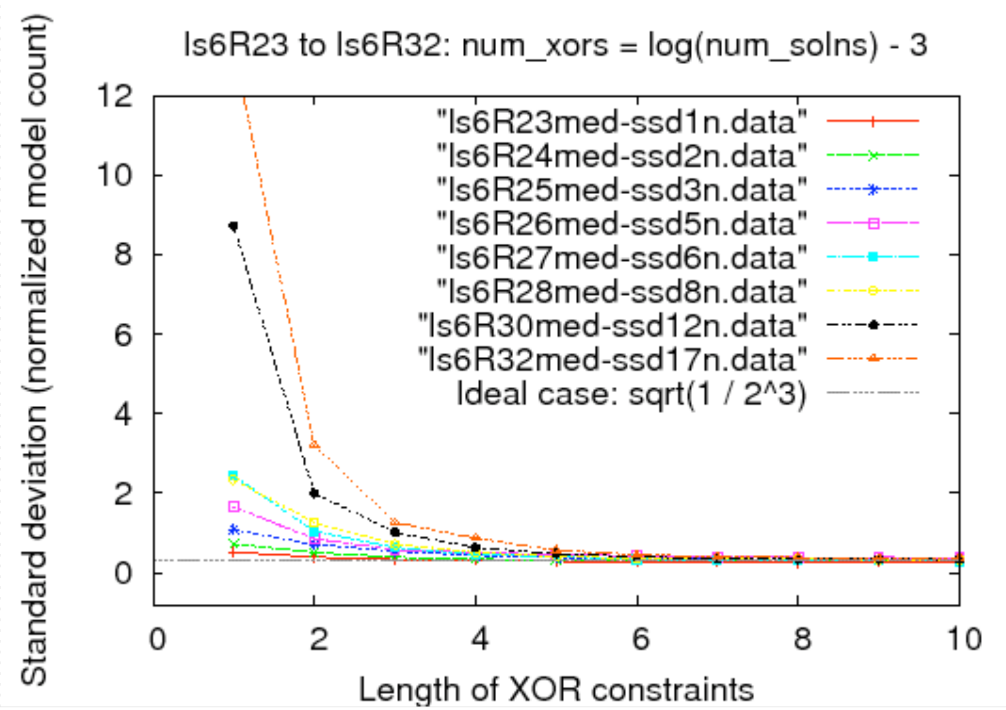

(quasi-group with holes)

80-140 variables
2^4-2^{20} solutions

3 remaining XORs

Ideal XOR length: 40-70

- Sample standard deviation initially decreases rapidly
- XOR lengths 5-7: already quite close to the ideal curve
- Not much change in s.s.d. after a while
  - Medium size XORs don't pay off well

Latin Square Formulas, Order 7

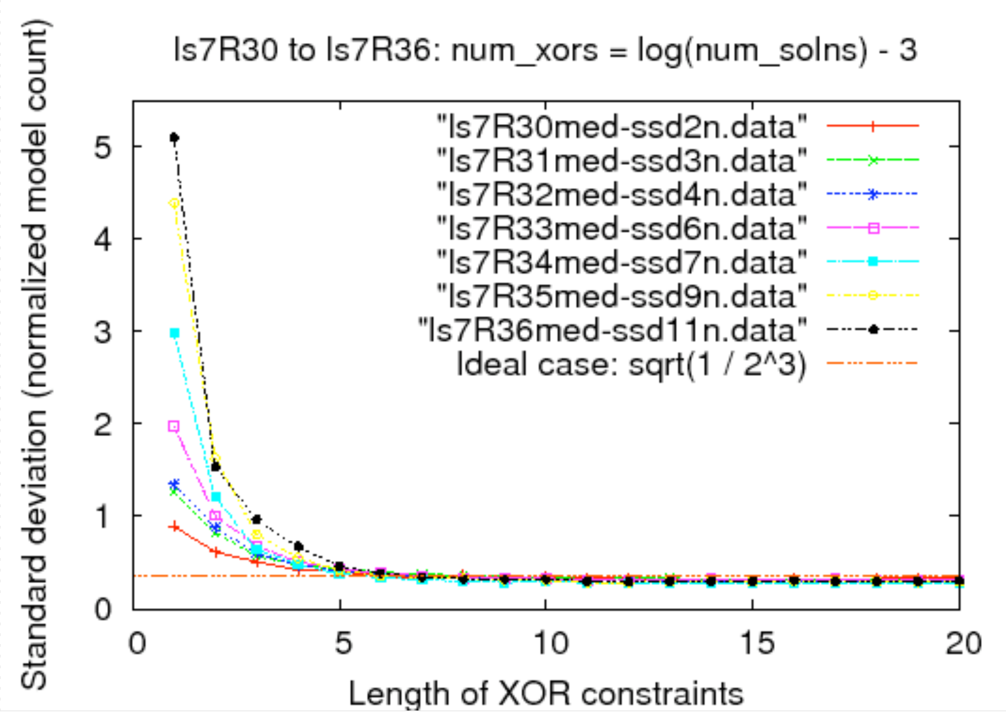

100-150 variables
2^5-2^{14} solutions

3 remaining XORs

Ideal XOR length: 50-75

- Similar behavior as Latin sq. of order 6
  - Even lower variance!
- Instances with more solutions have larger variance

Logistics Planning Instance

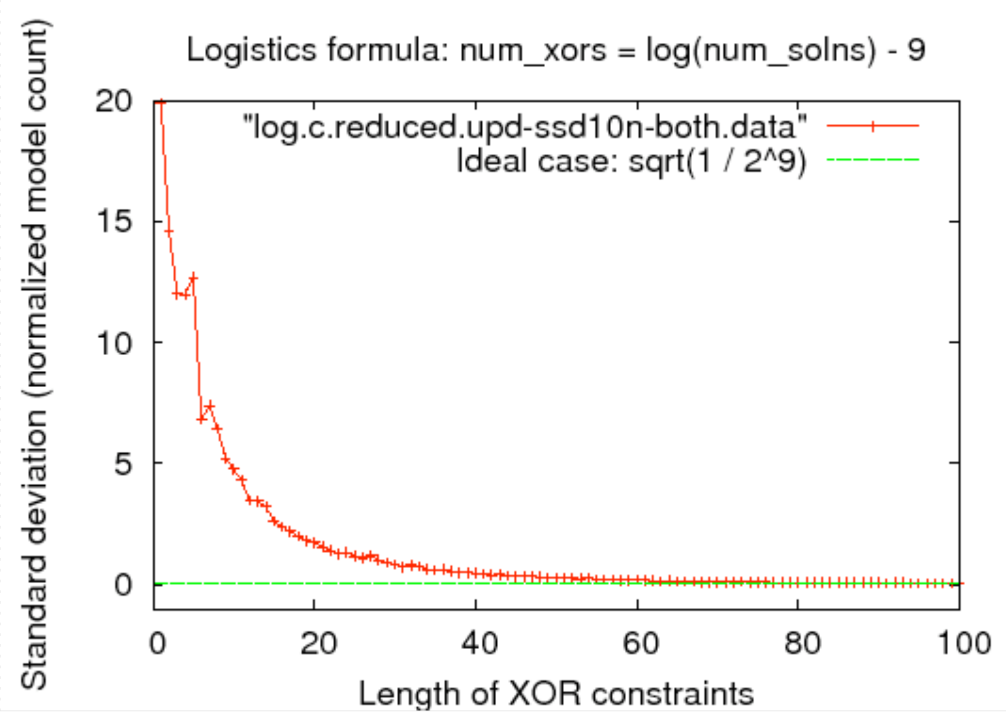

352 variables
2^{19} solutions

9 remaining XORs

Ideal XOR length: 151

- Variance drops sharply till XOR length 25
- XOR lengths 40-50 : very close to ideal behavior

Circuit Synthesis Problem

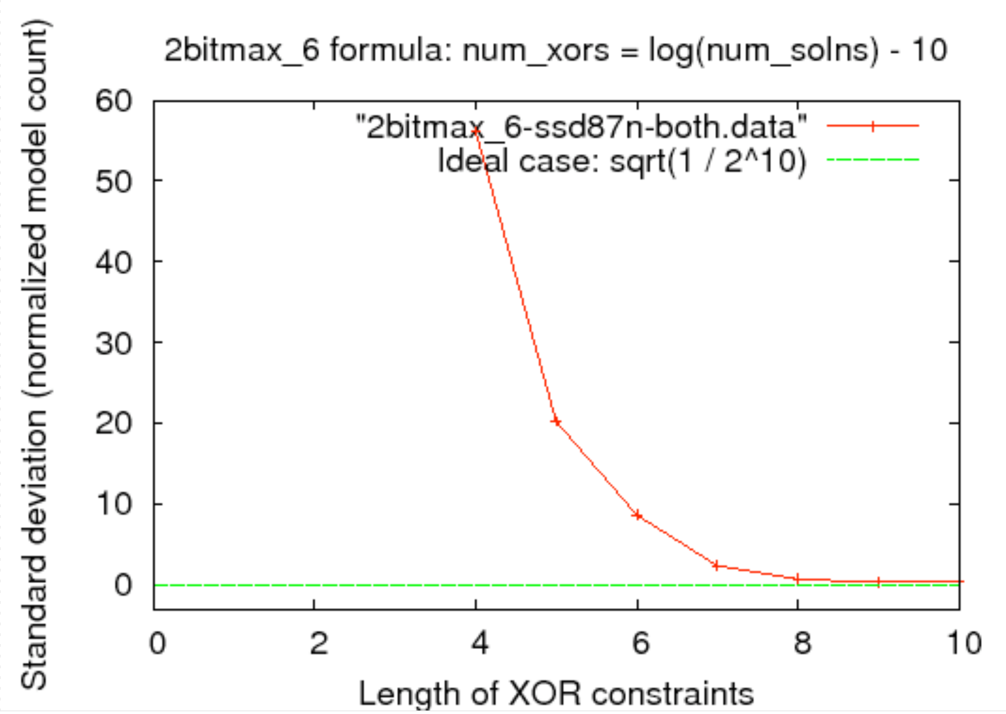

252 variables
2^{97} solutions

10 remaining XORs

Ideal XOR length: 126

- Standard deviation quite high initially
- But drops dramatically till length 7-8
- Quite close to ideal curve at length 10

Random 3-CNF Formulas

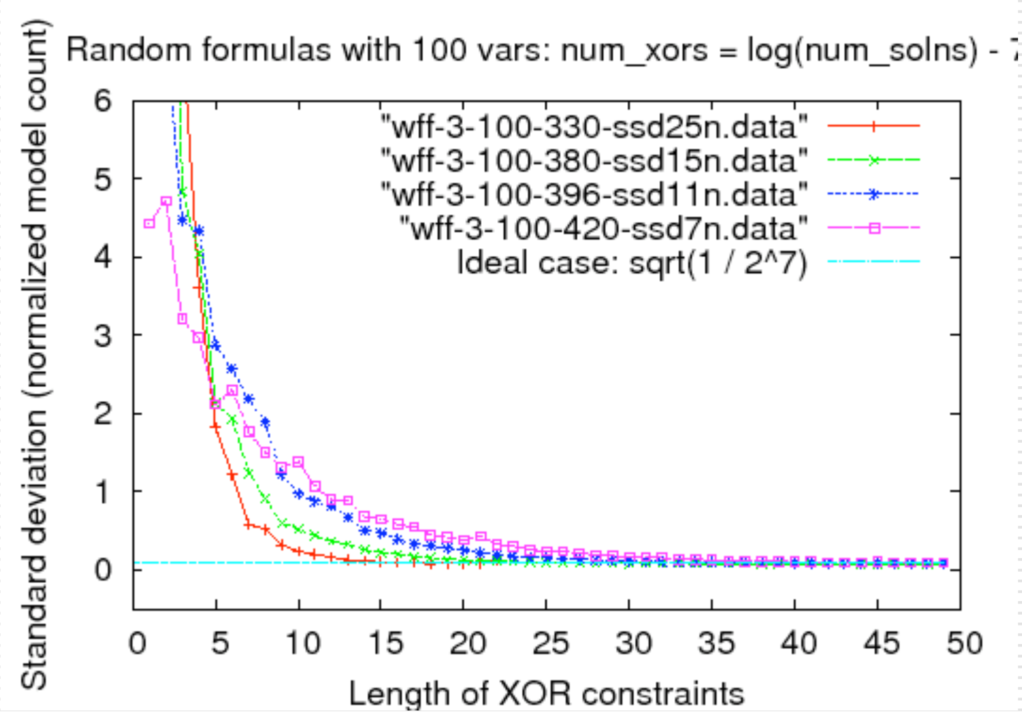

100 variables

2^{32}-2^{14} solutions

7 remaining XORs

Ideal XOR length: 50

- XORs don't behave as good as in structured instances
  - E.g. formulas at ratio 4.2 needs length 40+ (ideal: 50)
- Surprisingly, short XORs better at lower ratios!
  - Recall: model counting observed to be harder at lower ratios

Understanding Short XORs

What is it that makes short XORs work / not work well?

Backbone of the solutions provides some insight.

Intuitively,

- Large backbone
  - ⇒ short XOR often involves only backbone variables
  - ⇒ all or no solutions survive
  - ⇒ high variance

- Small backbone or split (local) backbones
  - ⇒ XOR involves non-backbone variables
  - ⇒ some solutions survive no matter what
  - ⇒ lower variance

Fixed-Backbone Formulas

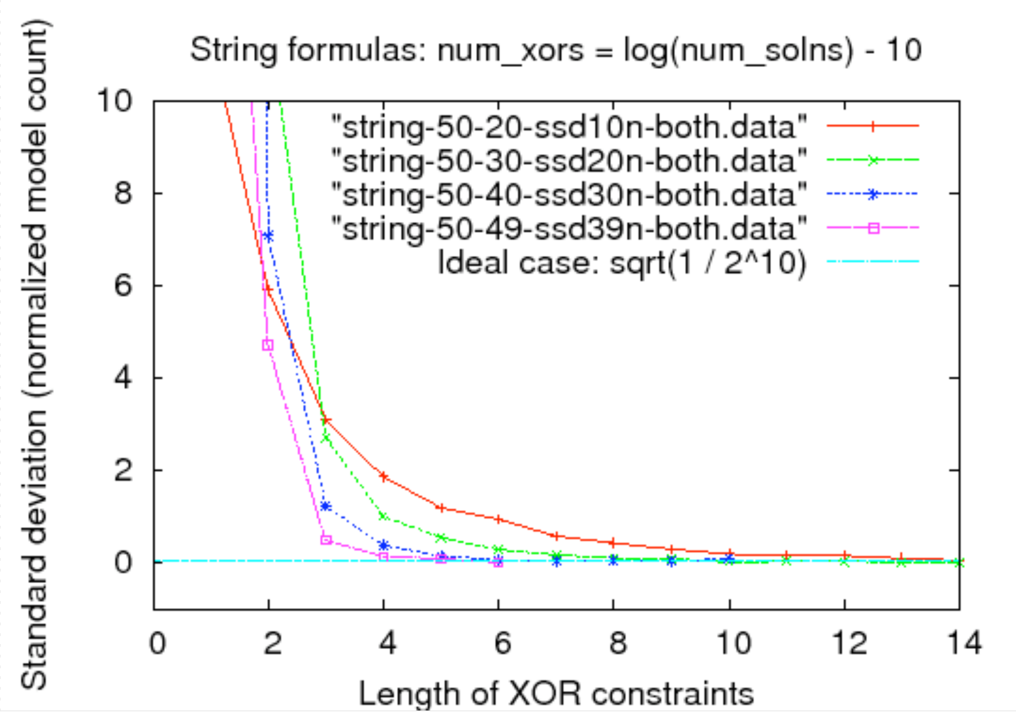

50 variables
2^{20}-2^{49} solutions

10 remaining XORs

Ideal XOR length: 25

- As backbone size decreases, shorter and shorter XORs begin to perform well

Interleaved-Backbone Formulas

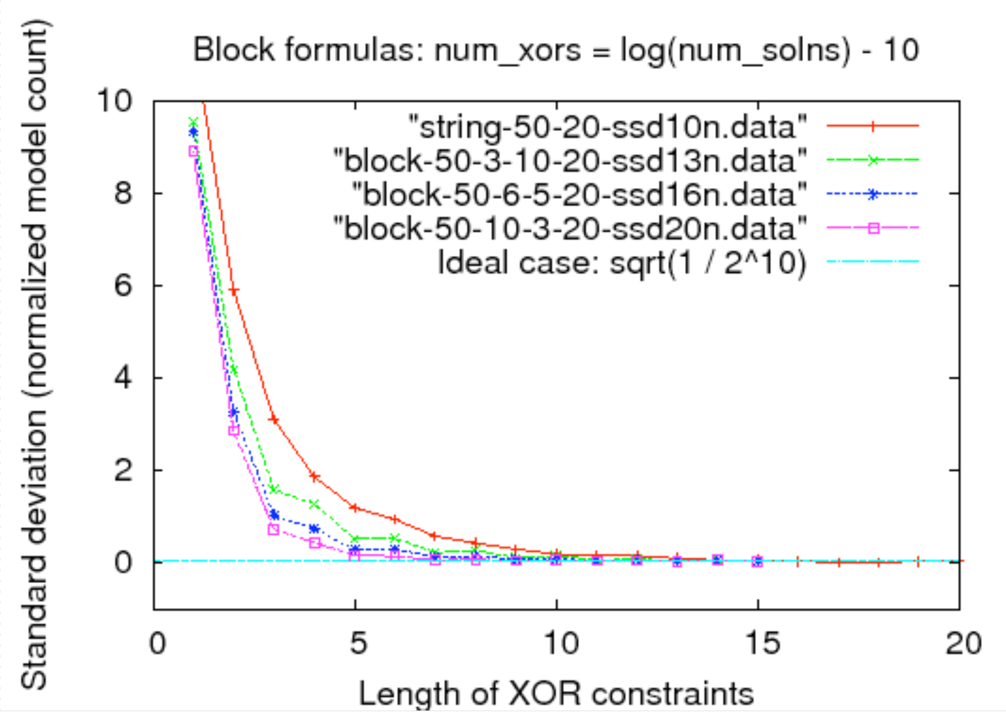

50 variables
2^{20}-2^{40} solutions

10 remaining XORs

Ideal XOR length: 25

- As backbone is split into more and more interleaved clusters backbones, shorter and shorter XORs begin to work well

Summary

- Short XORs can perform surprising well in practice for model counting and sampling
- Variance reduces dramatically at low XOR lengths
  - Increasing XOR length pays off quite well initially but not so much later
- Variance relates to solution backbones

Slides available at the SAT-07 poster session on Thursday!