Technology Overview

Modern SAT Solvers: Key Advances and Applications

Ashish Sabharwal
IBM Watson Research Center

Feb 23, 2011

Thanks…

For sharing their valuable knowledge, wisdom, slides, and more:

- Paul Beame, Henry Kautz (1998-2005, UW)
- Carla Gomes, Bart Selman (2005-2010, Cornell)
- Too many other colleagues to begin to mention here…

For exciting opportunities that lie ahead:

- "Watson" and colleagues (2010-, IBM Watson)

For this wonderful chance to be back at UW!

- Mausam

Outline

- The challenge of automated reasoning
- Dramatic progress in SAT solvers in the last 20 years

- Applications of SAT

- Key advances in SAT reasoning technology
- Advances in our understanding

- Looking ahead: *where is SAT research headed?*
- Summary

The Quest for Machine Reasoning

A cornerstone of Artificial Intelligence

Objective: Develop foundations and technology to enable effective, practical, large-scale automated reasoning

Machine Reasoning (1960-90s) → **Current reasoning technology**

Computational complexity of reasoning *appeared* to severely limit real-world applications

Revisiting the challenge: Significant progress with new ideas / tools for dealing with **complexity (scale-up), uncertainty, and multi-agent reasoning**

General Automated Reasoning

Problem instance → Model Generator (Encoder) → General Inference Engine → Solution

Domain-specific
e.g. logistics, chess, planning, verification, ...

Generic
applicable to all domains within range of modeling language

<u>Research objective</u>

Better reasoning and modeling technology

<u>Impact</u>

Faster solutions in several domains

Challenges in Automated Reasoning

Scalability

overcome high worst-case complexity

Robustness

guarantees against catastrophic failures

Synergy between paradigms

exact vs. probabilistic inference

Representation and balance

exploration vs. inference

Multi-agent reasoning

strategic decision making against adversary

Intelligent assistance

assist humans by solving sub-problems

Exponential Complexity Growth:
The Challenge of Complex Domains

Note: *rough estimates, for propositional reasoning*

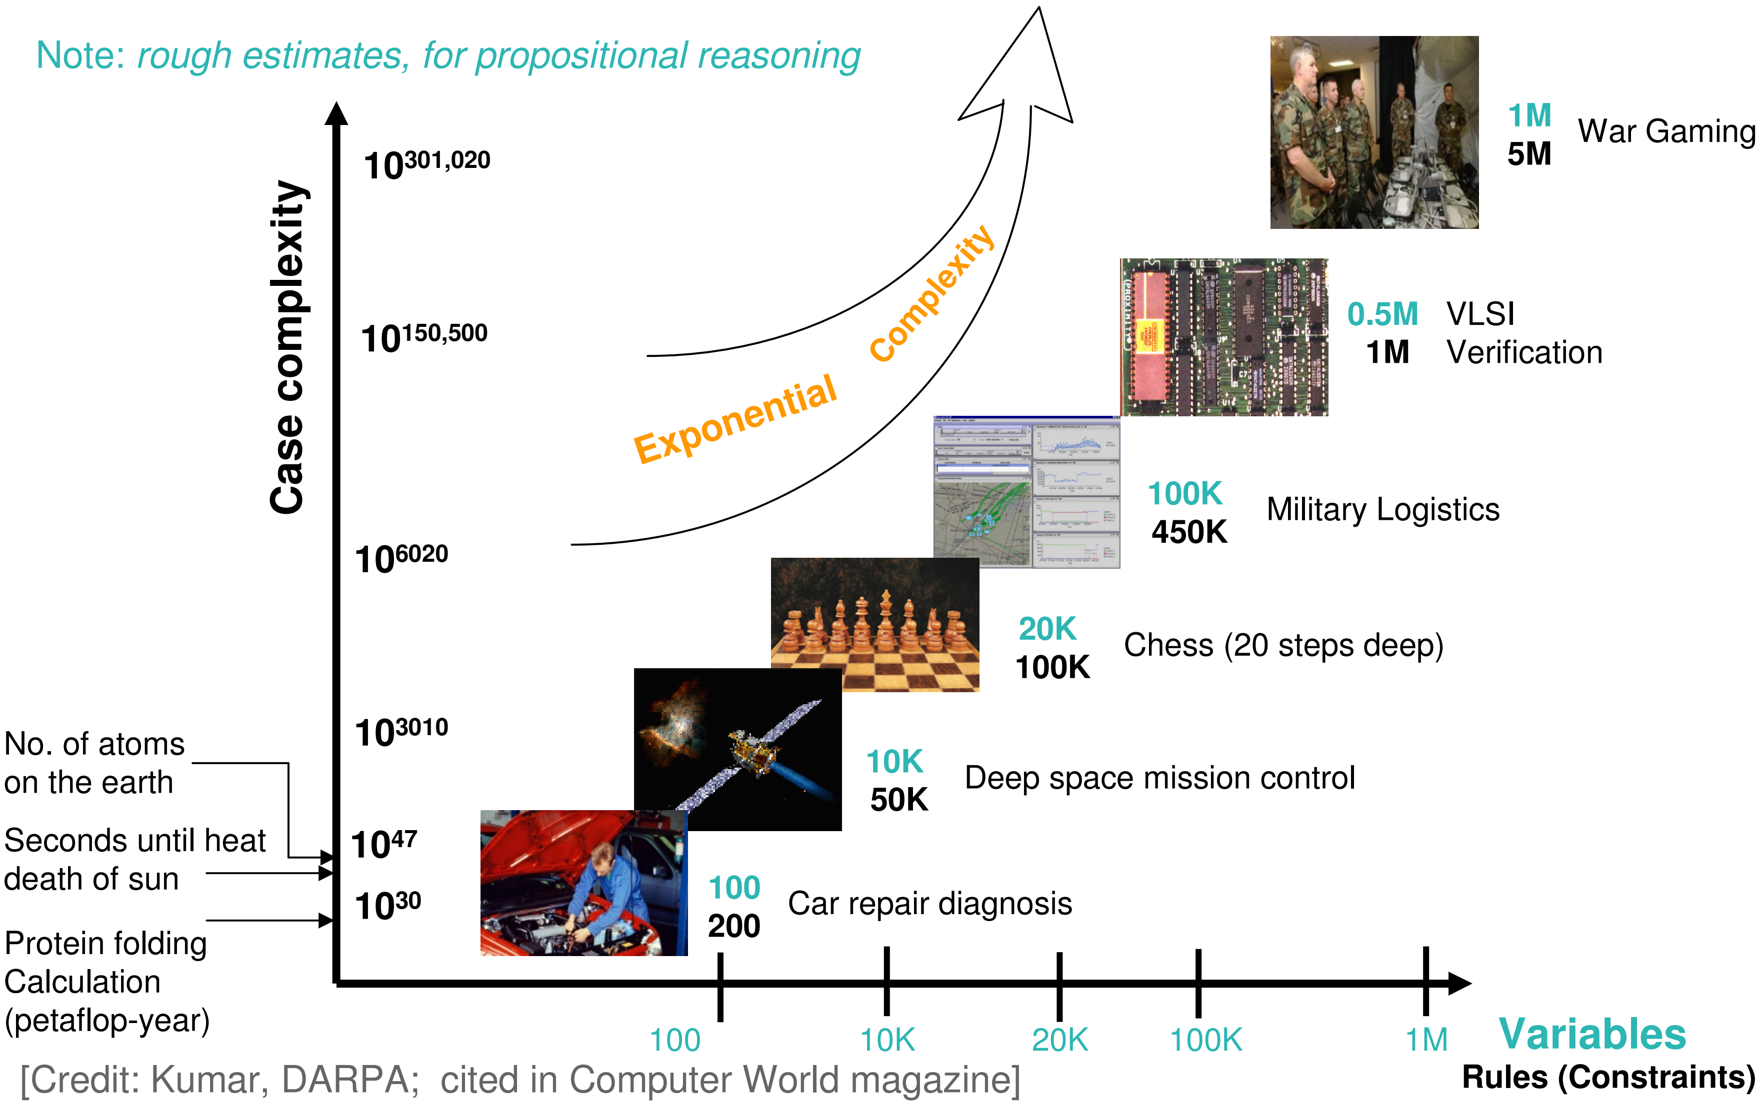

[Credit: Kumar, DARPA; cited in Computer World magazine]

Outline

- ✓ The challenge of automated reasoning
- Dramatic progress in SAT solvers in the last 20 years

- Applications of SAT

- Key advances in SAT reasoning technology
- Advances in our understanding

- Looking ahead: *where is SAT research headed?*
- Summary

Progress in the Last 20 Years

Focus: Combinatorial Search Spaces
specifically, the Boolean satisfiability problem, SAT

Significant progress since the 1990's. How much?

- Problem size: **We went from 100 variables, 200 constraints (early 90's) to 1,000,000+ variables and 5,000,000+ constraints in 20 years**

  Search space: from 10^30 to 10^300,000.
  [Aside: "one can encode quite a bit in 1M variables."]

- Is this just Moore's Law? It helped, but not much…
  - 2x faster computers does *not* mean can solve 2x larger instances
  - search difficulty does not scale linearly with problem size!
- **Tools**: 50+ competitive SAT solvers available

How Large are the Problems?

A bounded model checking (BMC) problem:

From "SATLIB":

http://www.satlib.org/benchm.html

SAT-encoded bounded model checking instances (contributed by Ofer Shtrichman)

In Bounded Model Checking (BMC) [BCCZ99], a rather newly introduced problem in formal methods, the task is to check whether a given model M (typically a hardware design) satisfies a temporal property P in all paths with length less or equal to some bound k. The BMC problem can be efficiently reduced to a propositional satisfiability problem, and in fact if the property is in the form of an invariant (Invariants are the most common type of properties, and many other temporal properties can be reduced to their form. It has the form of 'it is always true that ... '.), it has a structure which is similar to many AI planning problems.

Source: IBM

SAT Encoding

(automatically generated from problem specification)

```
p cnf 51639 368352
−1 7 0
−1 6 0
−1 5 0
−1 −4 0
−1 3 0
−1 2 0
−1 −8 0
−9 15 0
−9 14 0
−9 13 0
−9 −12 0
−9 11 0
−9 10 0
−9 −16 0
−17 23 0
−17 22 0
```

i.e., ((not x_1) or x_7)
((not x_1) or x_6)
etc.

x_1, x_2, x_3, etc. are our Boolean variables (to be set to True or False)

Should x_1 be set to False??

10 Pages Later:

```
185 −9 0
185 −1 0
177 169 161 153 145 137 129 121 113 105 97
 89 81 73 65 57 49 41
 33 25 17 9 1 −185 0
186 −187 0
186 −188 0
  ...
```

**i.e., (x_{177} or x_{169} or x_{161} or x_{153} …
x_{33} or x_{25} or x_{17} or x_{9} or x_{1} or (not x_{185}))**

clauses / constraints are getting more interesting…

Note x_1 …

4,000 Pages Later:

10236 −10050 0
10236 −10051 0
10236 −10235 0
10008 10009 10010 10011 10012 10013 10014
10015 10016 10017 10018 10019 10020 10021
10022 10023 10024 10025 10026 10027 10028
10029 10030 10031 10032 10033 10034 10035
10036 10037 10086 10087 10088 10089 10090
10091 10092 10093 10094 10095 10096 10097
10098 10099 10100 10101 10102 10103 10104
10105 10106 10107 10108 −55 −54 53 −52 −51 50
10047 10048 10049 10050 10051 10235 −10236 0
10237 −10008 0
10237 −10009 0
10237 −10010 0

…

Finally, 15,000 Pages Later:

$-7\ 260\ 0$
$7\ -260\ 0$
$1072\ 1070\ 0$
$-15\ -14\ -13\ -12\ -11\ -10\ 0$
$-15\ -14\ -13\ -12\ -11\ 10\ 0$
$-15\ -14\ -13\ -12\ 11\ -10\ 0$
$-15\ -14\ -13\ -12\ 11\ 10\ 0$
$-7\ -6\ -5\ -4\ -3\ -2\ 0$
$-7\ -6\ -5\ -4\ -3\ 2\ 0$
$-7\ -6\ -5\ -4\ 3\ -2\ 0$
$-7\ -6\ -5\ -4\ 3\ 2\ 0$
$185\ 0$

Search space of truth assignments: $2^{50000} \approx 3.160699437 \cdot 10^{15051}$

Current SAT solvers solve this instance in just a few seconds!

SAT Solvers as Search Engines

- Systematic SAT solvers have become really efficient at searching fast

- E.g., on an IBM model checking instance from SAT Race 2006, with ~170k variables, 725k clauses, solvers such as MiniSat and RSat roughly:
  - Make 2000-5000 **decisions**/second
  - Deduce 600-1000 **conflicts**/second
  - **Learn** 600-1000 **clauses**/second (#clauses grows rapidly)
  - Restart every 1-2 seconds (aggressive restarts)

SAT Solver Progress Trend (1994 – 2001)

Solvers have continually improved over time

Instance	Posit' 94	Grasp' 96	Sato' 98	Chaff' 01
ssa2670-136	40.66s	1.20s	0.95s	0.02s
bf1355-638	1805.21s	0.11s	0.04s	0.01s
pret150_25	>3000s	0.21s	0.09s	0.01s
dubois100	>3000s	11.85s	0.08s	0.01s
aim200-2_0-no-1	>3000s	0.01s	< 0.01s	< 0.01s
2dlx_..._bug005	>3000s	>3000s	>3000s	2.90s
c6288	>3000s	>3000s	>3000s	>3000s

Source: MarquesSilva-02

Forces Driving Faster, Better SAT Solvers

From academically interesting to practically relevant

- **"Real" benchmarks**, with real interest in solving them
- Regular **SAT Solver Competitions** (Germany-89, Dimacs-93, China-96, SAT-02, SAT-03, …, SAT-07, SAT-09, SAT-2011)
  - "Industrial-instances-only" **SAT Races** (2008, 2010)
  - A tremendous resource! E.g., SAT Competition 2006 (Seattle):
  - 35+ solvers submitted, downloadable, mostly open source
  - 500+ industrial benchmarks, 1000+ other benchmarks

- 50,000+ benchmark instances available on the Internet

This constant improvement in SAT solvers is the key to the success of, e.g., SAT-based planning and verification

Coming up: SAT Competition 2011!

- www.satcompetition.org
  – To be organized along-side the SAT-2011 Conference, Ann Arbor, MI
  – Exciting event for the community
  – An annual ‘check’: *how much did we improve this year?*

  – My task as one of the judges:
    ensure fairness and competitive spirit go hand-in-hand!

SAT Competition 2011

A competitive event of the SAT 2011 Conference

June 19th - June 22nd 2011, Ann Arbor, MI, USA

- If you have a great idea to improve SAT solvers, start working on it!

  [submissions due March 2]

Outline

- ✓ The challenge of automated reasoning
- ✓ Dramatic progress in SAT solvers in the last 20 years

- Applications of SAT

- Key advances in SAT reasoning technology
- Advances in our understanding

- Looking ahead: *where is SAT research headed?*
- Summary

For details on the state of the art in SAT…

Handbook of Satisfiability

- Everything from historical perspectives, theoretical foundations, practical solvers, applications, extensions, …

[27 chapters contributed by numerous experts]

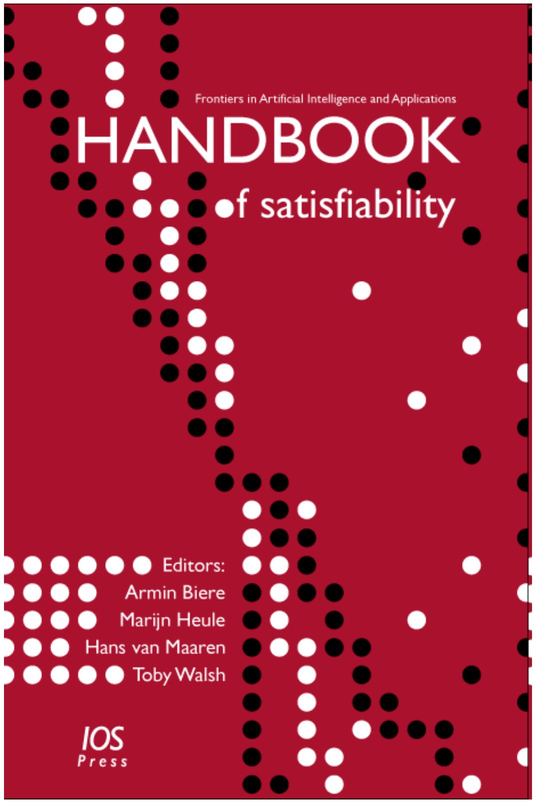

Applications of SAT (1)

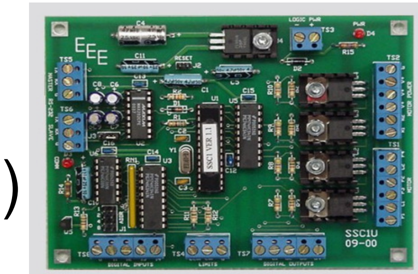

- Model Checking (hardware/software **verification**)
  - *Bounded* model checking (BMC)
  - Enhancement techniques: Abstraction and refinement
  - Major groups at Intel, IBM, Microsoft, and universities such as CMU, Cornell, and Princeton
  - **SAT has become a dominant back-end technology**
- Key questions:
  - **Safety** property: can system <u>ever</u> reach state S?
  - **Liveness** property: will system <u>always</u> reach T after it gets to S?
- Main idea:
  - Encode “reachability” in a *finite* transition graph
  - CNF encoding: in simple terms, “similar” to planning… next slide!

Applications of SAT (2)

- Classical Planning
  - Parallel step optimal plans
  - E.g., SATPLAN-06 fastest in this category at IPC-06 planning competition

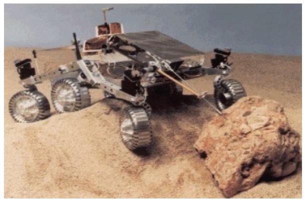

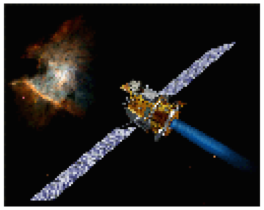

- Main idea:
  - Create planning graph, unfolded up to T steps
  - Encode legal transitions as CNF
    - **Variables**: action vars, fact vars *for each time step t*
    - **Clauses** encode action preconditions, effects, mutual exclusion, etc.
      e.g., (actionA_(t+1) ⇒ preA1_t AND preA2_t AND preA3_t),
      (NOT actionA_t OR NOT actionB_t), …

Applications of SAT (3)

- Combinatorial Design
  - Complex mathematical structures with desirable properties
    - Latin squares
    - Quasi groups
    - Ramsey numbers
    - Steiner systems
    - …
  - Highly useful for
    - design of experiments
    - coding theory, cryptography
    - drug design and testing
    - crop rotation schedules
    - …

Applications of SAT (4)

- Has been applied to solve *sub-problems* in many domains!

  – Test pattern generation
  – Scheduling
  – Optimal Control
  – Protocol Design (e.g., for networks)
  – Multi-agent systems
  – E-Commerce (E-auctions and electronic trading agents), etc.

Outline

- ✓ The challenge of automated reasoning
- ✓ Dramatic progress in SAT solvers in the last 20 years
- ✓ Applications of SAT
- Key advances in SAT reasoning technology
- Advances in our understanding
- Looking ahead: *where is SAT research headed?*
- Summary

Key Advances in SAT Reasoning

- Learn from mistakes *during* search
  - Clause Learning: when DPLL backtracks, learn a concise reason: what went wrong → avoid similar 'mistakes' in the future!
  - **Extremely powerful** in practice

Reason = a cut in implication graph

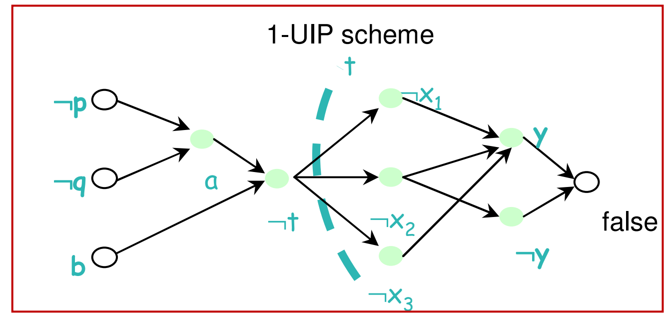

- Restart to avoid heavy-tailed behavior
  - SAT solvers known to sometimes "get stuck" and at other times be very fast… on similar problems, of similar size, etc.
  - Discovery: runtime distribution (across 'similar' instances or even on a single instance) is heavy-tailed → too many runs that can take way longer than the average!
  - **Solution**: restart computation differently, keep learned clauses!

Key Advances in SAT Reasoning

- Machine learning to build **algorithm portfolios**
  - Observation: no single SAT solver is good on every family of instances
  - Features of a given instance can be used to predict, with reasonable accuracy, which solver will work well on it!
  - **Solution**: design a portfolio solver using ML techniques
    - Based on runtime prediction models
    - Recent work – avoid complex models, use k-NN or clustering

- Automatic parameter tuning (generic and instance-specific)
  - SAT solvers are designed with many ‘hardwired’ parameters
  - Millions of parameter combinations – impossible to explore all by hand!
  - **Solution**: use automatic parameter tuning tools based on local search, genetic algorithms, etc.

Outline

- ✓ The challenge of automated reasoning
- ✓ Dramatic progress in SAT solvers in the last 20 years
- ✓ Applications of SAT
- ✓ Key advances in SAT reasoning technology
- Advances in our understanding
- Looking ahead: *where is SAT research headed?*
- Summary

Advances in Our Understanding (1)

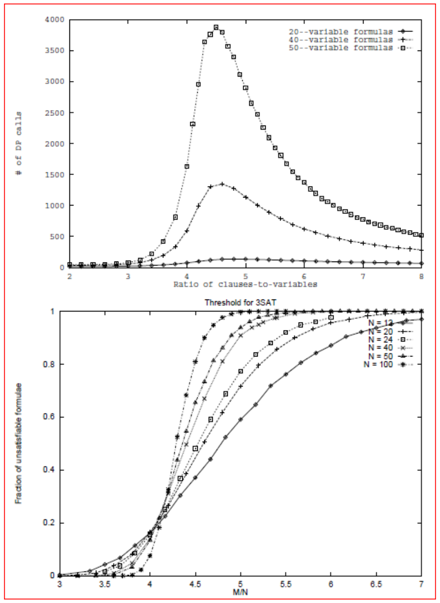

- Phase transitions and hardness profiles
  - Sharp threshold for satisfiability
  - Easy-hard-easy behavior
  - Heavy-tailed runtime phenomenon
    - Inspiring rapid restarts and portfolio methods

- Theoretical Foundations
  - SAT solvers as a “proof system” within Resolution
  - CL adds exponential power to basic DPLL
  - “CL + lots of restarts” as powerful as Resolution!

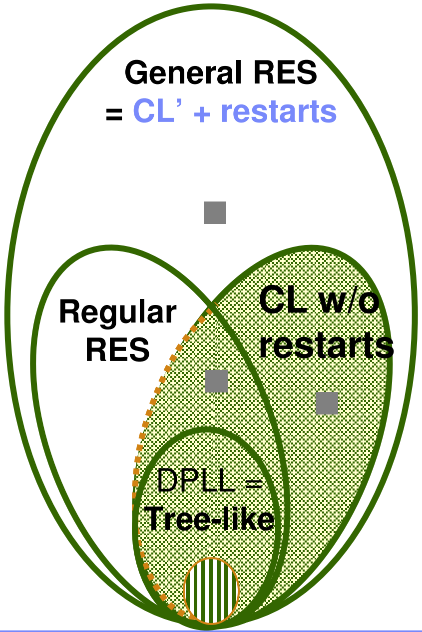

Advances in Our Understanding (2)

- Hidden Structure in “real-world” problems
  - **Backdoors** to satisfiability (and unsatisfiability)

constraint graph

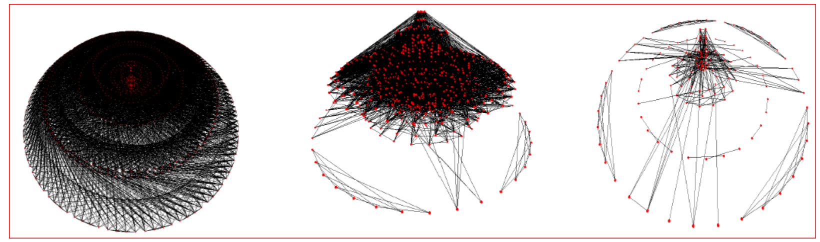

  - E.g., setting 16 variables (or even 5 vars) in this logistics planning problem with 843 vars, 7301 constraints simplifies it tremendously!
  - Note: such hidden structure *not* captured by traditional “cutsets”
- Beautiful insights into solution space: **clustering structure**!
  - Connections with Statistical Physics / **Survey Propagation**

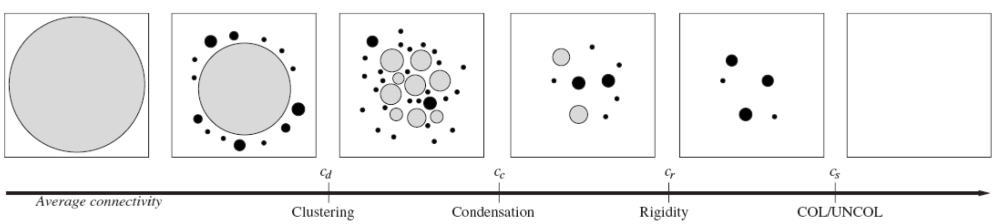

Outline

- ✓ The challenge of automated reasoning
- ✓ Dramatic progress in SAT solvers in the last 20 years
- ✓ Applications of SAT
- ✓ Key advances in SAT reasoning technology
- ✓ Advances in our understanding
- Looking ahead: *where is SAT research headed?*
- Summary

Looking Ahead: *Where is SAT Research Headed?*

- Direction A: getting more out of SAT solvers
  - Minimal/minimum **unsatisfiable cores**: very useful in practice!
  - **MAXSAT**, weighted MAXSAT
  - Circuit representations (rather than CNF)

- Direction B: tacking problems *harder* than SAT
  - Near-uniform **sampling** from the solution space
  - **Solution counting** (with relations to probabilistic inference)
    - #P-hard : challenging even to approximate with good confidence bounds

- Direction C: expanding the applicability of SAT technology
  - **Pseudo-Boolean** SAT (i.e., linear inequalities over Boolean vars)
  - **SMT**: Satifisiability Modulo Theories (e.g., linear arithmetic, bit-vector operations, uninterpreted functions)

Summary

- SAT research has been on an exciting 20 year journey
  – From toy problems with 100-200 variables
  – To industrial applications with 1M variables, 5M constraints

- Powerful new techniques and insights, many extensions, new connections with other areas of AI and beyond, strong groups in universities and industry alike

- Handbook of SAT: a great resource!

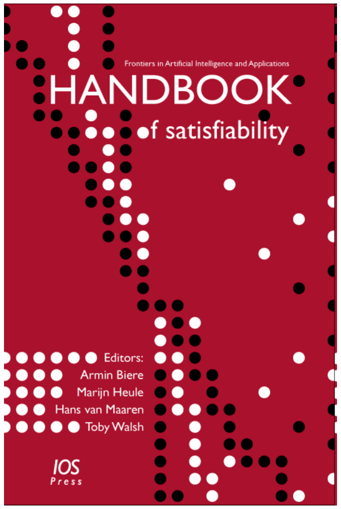

- Questions/comments/more information:
  ashish.sabharwal@us.ibm.com