

Research Interests

Ashish Sabharwal
Cornell University

The goal of my research is to *develop scalable and robust automated reasoning technology* that will allow computers to act intelligently in increasingly complex real-world settings and in competitive and uncertain environments. My research is driven by central applications and techniques in Artificial Intelligence (AI), and draws upon an extensive background in Theoretical Computer Science. I study combinatorial problems in domains such as planning, adversarial reasoning, verification and diagnosis, scheduling, design automation, resource economics, e-commerce, discrete mathematics, and, most recently, computational sustainability. My work has strong ties with probabilistic inference methods and the whole gamut of related applications.

A distinguishing strength of my research approach is *a unique blend of rigorous analysis and principled experimentation*, which I believe are equally critical components of fundamental research leading to practically useful technology. I place a strong emphasis on studying a variety of challenging combinatorial problems with researchers from diverse backgrounds and with varying interests. This has resulted in a series of fruitful collaborations in the past few years with theoreticians and practitioners alike.¹

Automated inference engines lie at the core of artificial intelligence, and are often embedded in systems ranging from car diagnosis tools to production planning software in the industry to complex control units in space missions. The vision of this technology is well exemplified by the DARPA Grand Challenge, where fully autonomous ground vehicles have relied on automated inference and control methods to successfully navigate over 100 miles of terrain without any human assistance. In 2007, the Urban Challenge took the setting from a desert course to an urban environment, safely navigating which revealed both the challenges in automated reasoning and decision making, as well as their tremendous potential in pushing the boundaries of computer science.

Combinatorial reasoning technology, in particular for propositional satisfiability (SAT), has witnessed an unprecedented growth. Starting with only a few hundred variable problems in the early 1990's, publicly available SAT solvers can now easily solve many large industrial problems with over 1 million variables and 5 million constraints. This research area promises to grow further in richness and impact as we begin to venture beyond basic combinatorial search, to problems such as *inferring properties of complex combinatorial spaces, and reasoning in multi-agent competitive environments*. This requires a shift in technology from NP-complete search problems to significantly more difficult #P- and PSPACE-complete problems. Efficient techniques for these problems are crucial for pushing automated inference engines to a level where they can address complex real-world systems. This poses many fundamental research challenges:

- *Scalability*: How can we overcome the high worst-case computational cost of solving these problems and scale them to the range of millions of variables demanded by real-world applications?
- *Robustness*: How can we obtain performance guarantees and have solutions that are robust against rare but catastrophic failures, such as power grids failures and stock markets crashes?
- *Synergy between exact and probabilistic inference*: How can we directly or indirectly use techniques developed for exact inference to solve problems in probabilistic inference, and vice versa?
- *Multi-agent reasoning*: How can we perform strategic decision making that will provably work in the presence of multiple agents and adversaries with competing interests?
- *Intelligent assistance*: How can the ability of computers to do fast exact inference on sub-problems be used to fruitfully assist humans in performing more complex tasks?
- *Representation and Balance*: How rich should the problem representation language be? How should we balance brute-force exploration vs. inference at each step of the search?

I believe addressing these issues will require a constant flow of ideas between foundational research and system design, and successful solutions will pave the road to a whole new range of applications.

¹The work discussed here was a joint effort with Carlos Ansotegui, Paul Beame, Jon Conrad, Carmel Domshlak, Carla Gomes, Willem-Jan van Hoeve, Jörg Hoffmann, Russell Impagliazzo, Henry Kautz, Andrew McCallum, Gilles Pesant, Toniann Pitassi, Ron Raz, Louis-Martin Rousseau, Meinolf Sellmann, and Bart Selman; and (then) students: Josh Buresh-Oppenheim, Matthew Cary, Aron Culotta, Bistra Dilkina, Justin Hart, Lukas Kroc, Yuri Malitski, Atri Rudra, Jordan Suter, and Erik Vee.

Main Research Themes

My research interests and methodology fall into three broad themes, which I discuss next along with a synopsis of my past work aligning with each theme.

Developing Scalable Automated Reasoning Methods

A key focus of my work is on building fast practical methods for combinatorial reasoning, often based on non-traditional approaches. The application domains and annual competitions that have been the catalyst for the tremendous growth in SAT solver and related technology have also brought the discipline to the boundary between general scientific research and careful, detailed engineering to excel on specific benchmark problems. A substantial amount of work goes into fine-tuning and exploring variations of techniques already in place. While this undoubtedly has its own merit, it is also crucial to look at problems from a fresh perspective. *Can one address issues regularly encountered by researchers by introducing a fundamentally new way of solving, or even representing, the problem?* This has served as a constant motivating question for my research. Here are three examples of innovative approaches that I have introduced, all of which have pushed the limits of automated reasoning by orders of magnitude:

XOR-STREAMLINING FOR MODEL COUNTING AND SAMPLING [AAAI-06/07, NIPS-06, IJCAI-07, SAT-07, ISAIM-08]: Implemented in a series of state-of-the-art tools (`MBound`, `SampleCount`, and `XorSample`) for counting and near-uniformly sampling solutions of combinatorial problems, XOR-streamlining is a fundamentally different technique for attacking these #P-hard problems than the traditional ones based on DPLL-search and Markov Chain Monte Carlo (MCMC) methods. Inspired by work in complexity theory, this approach introduced the first effective and scalable method for using a complete or local search SAT-solver essentially *off-the-shelf*, and a new framework for obtaining bounds with probabilistic correctness guarantees. This work was recognized with an Outstanding Paper Award at the 21st National Conference on Artificial Intelligence, AAAI-06. My recent work broadens the scope of these ideas by incorporating message passing techniques from probabilistic inference (namely, Belief Propagation) as well as statistical estimation.

DUAL FORMULATION FOR QBF SOLVERS [SAT-07]: Implemented as the solver `Duaffle`, the method represents a departure from the commonly used CNF-based representation formalism for quantified Boolean formulas. By using a new dual CNF-DNF representation based on a two-player game perspective of QBF domains, this approach brings out the full power of DNF-based “solution learning” techniques and facilitates, for the first time, constraint propagation across quantifiers—a bottleneck for search-based QBF solvers.

SYMMETRY-BREAKING AND REASONING BEYOND RESOLUTION [AAAI-05, Constraints J. 08]: Implemented in the structure-aware SAT solver `SymChaff`, this technique is again a departure from commonly used methods like “symmetry-breaking constraints” for SAT and CSP. By retaining and exploiting automatically generated contextual information about problem variables, it can achieve as much as (provably) exponential speed-ups over the best alternatives. The theme is the same as in the QBF work above: overcome limitations introduced by traditional CNF-based encodings by altering the representation formalism.

Formal Analysis of Inference Techniques Used in Practice

Understanding inherent strengths and limitations of various methods used in practice plays a crucial role in further development and successful application of such methods. Following this philosophy, I have created formal frameworks for understanding and analyzing state-of-the-art approaches in satisfiability testing and AI planning, complementing empirical observations about when these approaches work well and when they don’t. This has led to new ideas addressing fundamental limitations of known techniques. My methodology here has combined a mathematical analysis, often based on proof complexity theory, and systematic experimentation. Examples of my work in this direction include:

RESOLUTION-BASED FORMAL FRAMEWORK FOR SAT SOLVERS [IJCAI-03, SAT-03, JAIR-04]: This work, recognized as the Runner-up for the IJCAI-JAIR Best Paper Prize for 2003-2008, introduced the first formal proof complexity based framework for a rigorous analysis of key techniques often engineered in DPLL-style SAT solver implementations. It revealed the inherent power of certain learning and restart techniques, by relating them to the ‘resolution’ proof system. This laid a much needed foundation, helping make these techniques an integral part of the next generation of solvers, and spawning further research and improvement.

PROBABILISTIC INFERENCE AND STATISTICAL PHYSICS [UAI-07, NIPS-08, CPAIOR-08, ISAIM-08, SAC-09, in progress] A relatively recent message passing approach originating in statistical physics, namely Survey Propagation (SP), has turned out to be much more efficient than mainstream DPLL and local search methods for solving very hard random SAT instances. This work provided key insights into this somewhat mysterious technique, revealing through extensive experimentation what probabilistic information about the solution space does Survey Propagation effectively compute. The work also demonstrated how such properties of the solution space can be fruitfully exploited to obtain the number of solutions of these problems, and has recently provided the first clear formal connection between SP and statistics over “clusters” of solutions. This “cluster-centered” approach appears to be very promising for designing a new class of search algorithms.

ABSTRACTION IN AI PLANNING [ICAPS-06, JAIR (under review)]: Abstraction is a commonly employed technique, especially in model checking and planning, for improving efficiency by abstracting away non-critical details. Can abstraction methods really achieve any benefit in AI planning systems? This work provided a rather surprising *negative* answer for the best-case behavior of Resolution-based planners. This showed that the “informedness” of the search method must compete with the informedness of the abstraction heuristic, providing new insights into the design of abstraction techniques.

HARDNESS PROFILES AND PROBLEM STRUCTURE [CPAIOR-07/08/09, CP-07, ISAIM-08]: CSPs often exhibit an intriguing pattern: an abrupt phase transition from being feasible to being infeasible as a key problem parameter is varied. This work, with direct application to “wildlife corridor” design for grizzly bears in the Northern Rockies, empirically revealed for the first time such phenomena—and a corresponding “easy-hard-easy” pattern—for problems that combine both constraint satisfaction and optimization aspects. In a different direction, my work has brought to light the fundamental strength of the notion of propagation-based “backdoor sets” used to characterize real-world structure in combinatorial problems and explain the astonishing scalability of SAT solvers on structured industrial benchmarks.

Addressing Foundational Issues: Algorithm Design and Proof Complexity

The third theme of my research involves addressing foundational issues underlying automated reasoning systems. Specifically, I design efficient algorithms for constraint solvers and characterize the strength of various “proof systems”. This has resulted in the first polynomial time algorithms in some cases, and NP-completeness or hardness of approximation results in other. Two examples of my work in this area are:

FILTERING ALGORITHMS FOR SPECIAL CONSTRAINTS [CP-06, Constraints J. 09]: This work introduced the first polynomial time filtering algorithm for a combinatorial constraint (the “sequence” constraint) that appears frequently in scheduling and design automation problems. This resolved a question that had been open for 10 years in the Constraint Programming (CP) community. This work was recognized with the Best Paper Award at the 12th International Conference on Principles and Practice of Constraint Programming, CP-06. In related work [ModRef-07, CPAIOR-08], our algorithms revealed exponential memory and runtime savings that higher-level set-based representations in CP can achieve.

RESOLUTION COMPLEXITY AND HARDNESS OF APPROXIMATION [Complexity-01, Comp. Complexity J.-07, FOCS-02, SIAM J. Computing-04]: This proof complexity work showed that almost all instances of some interesting co-NP complete graph problems require exponential size Resolution proofs of infeasibility, even to approximate within significant factors, thus providing a large family of structured formulas hard for the Resolution proof system. The methodology involved combinatorial and probabilistic analysis. The work also showed that a natural class of approximate optimization algorithms for these problems must fail on almost all instances. In related work, we proved that even stronger proof systems, such as bounded-depth Frege systems described in many logic texts, require exponential size proofs even for very weak pigeonhole formulas, strengthening previously known results in this area.

Other Interests

My broad interests have drawn me to many applications of combinatorial analysis spanning various areas in computer science, such as *game theory* [AAAI-07], *computational geometry* [J. CGTA (to appear)], *circuit complexity*, and *digital rights management* [MSR Internship]. Looking forward, I expect to continue this stimulating endeavor, and collaborate with researchers in these and other fields.