Theoretical Foundations of Applied SAT Solving (14w5101)
January 19-24, 2014

MEALS

*Breakfast (Buffet): 7:00–9:30, Sally Borden Building, Monday–Friday
*Lunch (Buffet): 11:30–13:30, Sally Borden Building, Monday–Friday
*Dinner (Buffet): 17:30–19:30 pm, Sally Borden Building, Sunday–Thursday
Coffee Breaks: As per daily schedule, in the foyer of the TransCanada Pipeline Pavilion (TCPL)
***Please remember to scan your meal card at the host/hostess station in the dining room for each meal.**

MEETING ROOMS

All lectures will be held in the lecture theater in the TransCanada Pipelines Pavilion (TCPL). An LCD projector, a laptop, a document camera, and blackboards are available for presentations.

SCHEDULE

Sunday

16:00	Check-in begins (Front Desk – Professional Development Centre – open 24 hours)
17:30–19:30	Buffet Dinner, Sally Borden Building
20:00–	Informal gathering in 2nd floor lounge, Corbett Hall Beverages and a small assortment of snacks are available on a cash honour system.

Monday

7:00–8:45	Breakfast
8:45–9:00	*Introduction and Welcome* by BIRS staff, TCPL
9:00–9:50	Marijn Heule *Mini-tutorial on conflict-driven clause learning (CDCL)*
9:55–10:45	Sam Buss *Mini-tutorial on proof complexity*
10:45–11:10	Coffee break
11:10–12:00	Matti Järvisalo *Mini-tutorial on preprocessing*
12:00–13:00	Lunch
13:00–14:00	Guided tour of The Banff Centre; meet in the 2nd floor lounge, Corbett Hall.
14:00–14:15	Group Photo; meet in foyer of TCPL (photograph will be taken outdoors so a jacket might be required).
14:15–14:35	Presentation of workshop participants
14:35–15:25	Jakob Nordström *Mini-tutorial on weak proof systems and connections to SAT solving*
15:25–16:00	Coffee break
16:00–16:25	Laurent Simon *Understanding the power of glue clauses*
16:30–17:20	Armin Biere *Where does SAT not work?*
17:30–19:30	Dinner
20:00–	Open problem session

Tuesday

7:00–9:00	Breakfast
9:00–9:50	Priyank Kalla *Leveraging Groebner bases and SAT for hardware/software verification*
9:55–10:45	Daniel Le Berre *Survey on integrating cutting planes in CDCL solvers*
10:45–11:10	Coffee break
11:10–12:00	Albert Atserias *Mini-tutorial on semialgebraic proof systems*
12:00–13:30	Lunch
13:30–15:00	Afternoon break
15:00–15:25	Allen Van Gelder *Elementary short refutations for the clique-coloring principle and for Tseitin odd-charge graph formulas in extended resolution*
15:25–16:00	Coffee break
16:00–16:25	Marijn Heule *Efficient and verified checking of unsatisfiability proofs*
16:30–17:20	Stefan Szeider *Survey of parameterized complexity and SAT*
17:30–19:30	Dinner

Wednesday

7:00–9:00	Breakfast
9:00–9:50	Joao Marques-Silva *Problem solving with SAT oracles*
9:55–10:45	Albert Oliveras *Survey of satisfiability modulo theories (SMT)*
10:45–11:10	Coffee break
11:10–11:35	Norbert Manthey *Recent developments in parallel SAT solving*
11:40–12:05	Ashish Sabharwal *Resolution and parallelizability: Barriers to the efficient parallelization of SAT solvers*
12:05–13:30	Lunch
	Free afternoon
17:30–19:30	Dinner

Thursday

7:00–9:00	Breakfast
9:00–9:50	Martina Seidl *Recent trends in QBF solving*
9:55–10:20	Nina Narodytska *Reactive synthesis via QBF solving*
10:20–10:45	Rahul Santhanam *Beating brute force search for QBF satisfiability*
10:45–11:10	Coffee break
11:10–11:35	Moshe Vardi *Phase transitions and computational complexity*
11:40–12:05	Sean Weaver *Satisfiability-based set membership filters*
12:05–13:30	Lunch
13:30–15:00	Afternoon break
15:00–15:25	Paul Beame *Exact model counting: SAT-solver based methods versus lifted inference*
15:25–16:00	Coffee break
16:00–16:25	Nicola Galesi *Space complexity in algebraic proof systems*
16:30–16:55	Massimo Lauria *Narrow proofs may be maximally long*
17:00–17:25	Jan Johannsen *Lower bounds for width-restricted clause learning*
17:30–19:30	Dinner
20:00–	Tentative second open problem session or discussion

Friday

7:00–9:00	Breakfast
9:00–9:25	Carsten Sinz *Abstraction and multi-encodings in SAT*
9:30–9:55	Oliver Kullmann *Unit-clause propagation and monotone circuits*
10:00–10:25	Chris Beck *Strong ETH holds for regular resolution*
10:25–11:00	Coffee break
11:00–11:25	Denis Bueno *Detecting traditional packing, decisively*
11:30–11:55	Vijay Ganesh *Timed PageRank and branching heuristics in CDCL SAT solvers*
12:00–13:30	Lunch

Friday departure information

- At 8:45–9:00, a bellman comes with a van and picks up participant luggage from the front of Corbett Hall for storage at the front desk.
- Check-out from the guest rooms is by 12:00.
- Workshop participants are welcome to use BIRS facilities (BIRS coffee lounge, TCPL and reading room) until 15:00.

Theoretical Foundations of Applied SAT Solving (14w5101)
January 19-24, 2014

ABSTRACTS

Speaker: **Albert Atserias** (Universitat Politecnica de Catalunya)
Title: *Mini-tutorial on semialgebraic proof systems*
Abstract:

A variety of semialgebraic proof systems, i.e. those operating with polynomial inequalities over the reals, were defined in the last two decades to reason about optimality or near optimality in combinatorial optimization problems. In this mini-tutorial we overview their origins and also compare them to the traditional proof systems for propositional logic. From a proof complexity point of view, the main advantage that semialgebraic proof systems offer over low-level logic-based systems is the greater expressive power of their proof lines, while preserving certain good algorithmic properties for the proof-search problem. These good algorithmic properties also make them amenable to analysis by the proof-complexity methods to prove lower bounds. We will try to illustrate these points through some examples from the literature.

Speaker: **Paul Beame** (University of Washington)
Title: *Exact model counting: SAT-solver based methods versus lifted inference*
Abstract:

The best current methods for exactly computing the number of satisfying assignments, or the satisfying probability, of Boolean formulas are based on SAT solvers enhanced with component caching. These can be seen, either directly or indirectly, as building decision-DNNF (decision decomposable negation normal form) representations of their input Boolean formulas. We show that such representations, and indeed those that employ a more general form of component decomposition, can be converted into read-once branching programs (ROBPs) with only a quasi-polynomial increase in size. We use this together with new exponential lower bounds on the ROBP size of some simple natural DNF formulas associated with queries considered in probabilistic databases to derive exponential lower bounds on the sizes of these representations and, therefore, on this approach to exact model counting, including queries for which there are simple polynomial algorithms to compute these counts using "lifted" inference methods.

Joint work with Jerry Li, Sudeepa Roy, and Dan Suciu.

Speaker: **Chris Beck** (Princeton University)
Title: *Strong ETH holds for regular resolution*
Abstract:

Let c_k be the infimum of real numbers so that k-SAT is solvable in time $2^{c_k n}\text{poly}(n)$. The Strong Exponential Time Hypothesis is that the limit of the sequence c_k is 1, i.e., that the difficulty of k-SAT approaches exhaustive search as k increases. This is true for the best algorithms currently known for k-SAT, as well as an empirical observation about the performance of SAT solvers on instances with large clauses.

Since many SAT solvers are based on the resolution proof system, lower bounds for this system give lower bounds for large families of such solvers. We show that no algorithm formalizable in the subsystem of regular resolution can be a counter-example to SETH. More precisely, we show that there are unsatisfiable k-CNF formulas on n variables so that any regular resolution refutation requires size $2^{1-\epsilon_k}n$, where $\epsilon_k = \tilde{O}(k^{1/4})$. We also improve the lower bounds for general resolution substantially, showing that the same formulas require general resolution size at least $(3/2)^{(1-\epsilon_k)n}$.

Speaker: **Armin Biere** (Johannes Kepler University)
Title: *Where does SAT not work?*
Abstract:

Even though SAT solving is quite successful in many application there are still open issues, where SAT does not work or at least we do not know how or why it works. In this talk we will go over some of these issues, including learning definitions through extended resolution, speeding up CDCL by local search effectively, arithmetic reasoning on the CNF level, data flow algorithms for SAT, limits of portfolio based parallel SAT solving, dynamic and more general reencoding, data structures for next generation SAT solvers, and last but not least the question of how and why VSIDS works.

Speaker: **Denis Bueno** (University of Michigan)
Title: *Detecting traditional packing, decisively*
Abstract:

Many important tasks in malicious software (malware) analysis rely on answering difficult questions: "Does this program exhibit a certain behavior? Is this code triggered to behave badly by some environmental condition?" It is widely believed that these questions are undecidable — that there exists no algorithm capable of answering "yes" or "no" correctly to all instances of these questions. As another concrete example, previous work has shown that detecting whether malware is "packed" — that it has a compressed or encrypted payload which is only revealed at runtime — is undecidable. Our work shows that under very broad and realistic assumptions about time and space resources, detecting packing is in fact NP-complete, rather than undecidable. Our talk will discuss the ideas that make our proofs possible.

Speaker: **Sam Buss** (University of California, San Diego)
Title: *Mini-tutorial on proof complexity*
Abstract:

This talk will survey propositional proof complexity. It will cover proof systems such as resolution, Frege systems, extended Frege systems, cutting planes, nullstellensatz and the polynomial calculus, with an emphasis on how they apply to SAT solvers or may possibly point the way to new algorithms for satisfiability. The talk will also discuss upper and lower bounds. It will discuss Craig interpolation, which has proved to be important both for lower bounds on propositional proofs and for applied SMT solvers.

Speaker: **Nicola Galesi** (Universita degli Studi di Roma La Sapienza)
Title: *Space complexity in algebraic proof systems*
Abstract:

The study of space measure in Proof Complexity has a gained in the last years more and more importance: it is clearly of theoretical importance in the study of complexity of proofs; it is connected with SAT solving, since it might provide theoretical explanations of efficiency or inefficiency of specific Theorem Provers or SAT solvers; it is connected with important characterizations studied in Finite Model Theory, thus providing a solid link between the two research fields.

We will present a recent work where we devise a new general combinatorial framework for proving space lower bounds in algebraic proof systems like Polynomial Calculus (PC) and Polynomial Calculus with Resolution (PCR). Our method can be view as a Spoiler-Duplicator game, which is capturing boolean reasoning on polynomials. Hence, for the first time, we move the problem of studying the space complexity for algebraic proof systems in the range of 2-players games, as is the case for Resolution. This can be seen as a first step towards a precise characterization of the space for algebraic systems in terms of combinatorial games, like Ehrenfeucht-Fraisse , which are used in Finite Model Theory.

A simple application of our method allows us to obtain all the currently known space lower bounds for PCR, like the Pigeonhole Principle. More importantly, using our approach in its full potentiality, we answer to the open problem of proving space lower bounds in Polynomial Calculus and Polynomials Calculus with Resolution for the polynomial encoding of randomly chosen k-CNF formulas. Our method also applies

to the the well studied Graph Pigeonhole Principle which is a Pigeonhole principle over a constant (left) degree bipartite expander graph.

The work arise several open problems might be solved generalizing our approach.

Speaker: **Vijay Ganesh** (University of Waterloo)
Title: *Timed PageRank and branching heuristics in CDCL SAT solvers*
Abstract:

While modern Conflict-driven Clause Learning SAT solvers employ a large number of techniques/data structures, there are 4 that play a crucial role in all of them, namely, conflict-driven clause learning with backjumping, VSIDS branching heuristic and variants, efficient implementation of boolean constant propagation, and restarts. In this talk, I will focus on the VSIDS branching heuristic which perhaps has been the least studied. We draw a parallel between the dynamic aspect of VSIDS with an online version of Timed PageRank (Eigenvector Centrality) ranking function that is used to determine the most influential node in a dynamically evolving graph. This parallel allows us to assert that Timed PageRank provides a good analytical model for VSIDS that is amenable to further mathematical analysis. We pose questions like "what problem is VSIDS trying to solve?", "why and how?", and attempt to answer them. The hope is that getting a better understanding of VSIDS will lead to much better SAT solvers.

Speaker: **Allen Van Gelder** (University of California Santa Cruz)
Title: *Elementary short refutations for the clique-coloring principle and for Tseitin odd-charge graph formulas in extended resolution*
Abstract:

The clique-coloring principle is a family of propositional CNF formulas that state, if an undirected graph G on n vertices has a clique of size s and $t < s$, then G cannot be t-colored. For $t = n - 1$ and $s = n$, this is the well known pigeon-hole principle, but the clique-coloring principle is of interest because some logics that have short (i.e., polynomial length) proofs of the pigeon-hole principle do not have short proofs of the clique-coloring principle (subject to some complexity assumptions in some cases). This paper describes an elementary extended resolution refutation for the clique-coloring principle. It uses Cook's well known short extended resolution refutation of the pigeon-hole principle. The constructions make heavy use of hyper-resolution as an organizing tool. Several recent papers attempt to incorporate extended resolution in disciplined ways into deterministic CDCL solvers, but none of these proposals consider the issue of whether hyper-resolution opportunities are thus created. Therefore one purpose of presenting the construction is to show the kind of extended proof that has proved to be much shorter than any non-extended proof. A second purpose is to show the kind of constructions that need to be avoided if a super-polynomial family for extended resolution is to be demonstrated.

A quite different family of hard formulas for resolution is the Tseitin odd-charge graph family. This paper describes an elementary extended resolution refutation for this family that brings its refutations down to polynomial length. Such an explicit construction seems to be unreported, but is not surprising. The question is what we can learn about the practical aspects of extended resolution by studying such constructions. Sat Modulo Theories (SMT) is suggested as a promising way forward.

Speaker: **Marijn Heule** (University of Texas at Austin)
Title: *Mini-tutorial on conflict-driven clause learning (CDCL)*
Abstract:

Satisfiability solvers have become powerful search engines to solve a wide range of applications in fields such as formal verification, planning and bio-informatics. Due to the elementary representation of SAT problems, many low-level optimizations can be implemented. At the same time, there exist clause-based techniques that can simulate several high-level reasoning methods. The tutorial focuses on the search procedures in the successful conflict-driven clause learning (CDCL) solvers. It explains how to learn from conflicts and provides an overview of effective heuristics for variable and value selection.

Speaker: **Marijn Heule** (University of Texas at Austin)
Title: *Efficient and verified checking of unsatisfiability proofs*
Abstract:

Satisfiability (SAT) solvers act as the core search engine in many tools used for bounded model checking and the verification of hardware and software. It is incumbent upon these solvers to produce correct results. For many tools, such as theorem provers, trusting the results of SAT solvers is not enough; they should be verified. Two problems obstruct wide-spread verification of unsatisfiability results: checking proofs can be very expensive and existing checking procedures cannot verify all existing techniques used in modern SAT solvers. We tackle both problems by introducing a new clausal proof format that allows efficient checking of all existing SAT solving techniques. Our proposed method has one disadvantage: a fast checker for the new format is more complex as compared to checkers for existing formats. We deal with this issue by creating a proof compressor that reduces the size of unsatisfiability proofs significantly. The compact proof can be validated using a simple verified checker. We implemented and mechanically verified such a proof checker using the ACL2 theorem prover.

Speaker: **Matti Järvisalo** (University of Helsinki)
Title: *Mini-tutorial on preprocessing*
Abstract:

The current standard approach to SAT-based problem solving consists of three main steps: (i) encoding (into propositional logic, finally to CNF); (ii) preprocessing (polynomial-time satisfiability-preserving rewriting of the CNF), and (iii) search (for satisfiability).

Preprocessing is today an integral part of this SAT solving workflow, significantly speeding up the actual search for satisfiability. Moreover, interleaving preprocessing steps with the actual CDCL search, referred to as 'inprocessing SAT solving', has recently proven to be a fruitful direction for further improving the efficiency of CDCL SAT solvers in practice.

This tutorial aims at giving a general overview of some of the currently most important SAT preprocessing techniques. As time permits, we will also overview formal underpinnings of inprocessing SAT solvers, and discuss some of the (dis)connections between theory and practice in view of preprocessing.

Speaker: **Jan Johannsen** (LMU Munich)
Title: *Lower bounds for width-restricted clause learning*
Abstract:

We show lower bounds for clause learning SAT algorithms (without restarts) when the width of learned clauses is restricted. These algorithms require time $2^{\Omega(n \log n)}$ to refute the pigeonhole principle clauses PHP_n when learning only clauses of width $n/2$. They also requite time $2^{\Omega(n)}$ to refute the ordering principle clauses Ord_n when learning only clauses of width $n/4$. In general, for unsatisfiable input formulas of small width, lower bounds for width-restricted clause learning followq from resolution lower bounds. All lower bounds obtained are of the same order of magnitude as known lower bounds for DPLL algorithms without clause learning.

Parts of this presentation are based on joint work with Sam Buss, Jan Hoffmann and Eli Ben-Sasson.

Speaker: **Priyank Kalla** (University of Utah)
Title: *Leveraging Groebner bases and SAT for hardware/software verification*
Abstract:

Advancements in SAT and SMT techniques have had a tremendous impact in electronic design automation, especially in hardware verification. SAT solvers, however, show particularly poor performance on verification of custom-designed arithmetic circuits, such as those found in signal processing, cryptography, error-correction coding, among others. For such applications, techniques from commutative algebra and algebraic geometry — i.e. modern Groebner bases theory and technology — prove to be more efficient. Conversely, Groebner basis techniques often fail on problems that are quickly solved by SAT.

SAT and Groebner basis techniques are therefore rather complementary, but often utilized separately. Unlike SAT, modern algebraic geometry does not "search" for the solutions; rather it reasons about dimension, presence/absence, intersection of the solution-sets. We therefore ask: can these different approaches be combined or adapted together in some manner? If so, what applications would benefit? In this talk, I will present a tutorial on Groebner bases theory, and review contemporary literature to show how Groebner bases techniques have been applied to SAT problems. Finally, I will describe some investigations by my research group that highlight the potential and challenges of integrating Groebner bases techniques with SAT search, particularly for hardware verification applications.

Speaker: **Oliver Kullmann** (Swansea University)
Title: *Unit-clause propagation and monotone circuits*
Abstract:

We show ([1]) that computation via unit-clause propagation and via monotone circuits are closely related, exploiting the basic ideas from [2].

We apply this to the question of finding "good" CNF-representations of systems of XOR-constraints (or "parity constraints"), in combination with the hardness-measurement approach as developed in [3]. The lower bound on the monotone circuit complexity of monotone span programs from [4] yields a super-polynomial size bound on "good" CNF-representations.

After this fundamental negative result, we turn to positive results for good CNF-representations on XOR's. Various hardness measures are considered.

References:

[1] Matthew Gwynne and Oliver Kullmann On SAT representations of XOR constraints, December 2013, http://arxiv.org/abs/1309.3060, conference version (to appear LATA 2014): On SAT representations of XOR constraints, http://www.cs.swan.ac.uk/ csoliver/papers.html#XOR2013LATA

[2] Christian Bessiere, George Katsirelos, Nina Narodytska, and Toby Walsh. Circuit complexity and decompositions of global constraints. In Twenty-First International Joint Conference on Artificial Intelligence (IJCAI-09), pages 412–418, 2009. http://arxiv.org/abs/0905.3757

[3] Matthew Gwynne and Oliver Kullmann Generalising unit-refutation completeness and SLUR via nested input resolution Journal of Automated Reasoning, January 2014, Volume 52, Issue 1, pages 31-65. http://cs.swan.ac.uk/ csoliver/papers.html#SLUR2013

[4] Laszlo Babai, Anna Gal, and Avi Wigderson. Superpolynomial lower bounds for monotone span programs. Combinatorica, 19(3):301319, March 1999. http://www.cs.utexas.edu/ panni/bgw.ps

Speaker: **Massimo Lauria** (KTH Royal Institute Of Technology)
Title: *Narrow proofs may be maximally long*
Abstract:

If a CNF has a resolution refutation of width w we can find it with in time $n^{O(w)}$, by inferring all possible clauses in width at most w. It is natural to ask whether this method can be improved.

We give a negative answer by showing that there are 3-CNF formulas over n variables, refutable in resolution in width w, but that require resolution proofs of size $n^{\Omega(w)}$.

Moreover, our lower bound extends to the more powerful proof systems polynomial calculus resolution (PCR) and Sherali-Adams, implying that the corresponding size upper bounds in terms of degree and rank are tight as well. In contrast, the formulas have Lasserre proofs of constant rank and size polynomial in both n and w.

Joint work with Albert Atserias and Jakob Nordström.

Speaker: **Daniel Le Berre** (Universite d'Artois)
Title: *Survey on integrating cutting planes in CDCL solvers*
Abstract:

CDCL solvers are based on the resolution proof system. It is tempting to extend such architecture to more powerful proof systems (e.g. cutting planes). There have been several attempts in that direction in

the past decade. We will review some of them in this talk. We will especially focus on the assumptions underlying those approaches, how they are closely related to the CDCL framework. We will define precisely what is the input of those solvers and which rules they use. Finally, we will study the proof generated by such an "extended CDCL solver" on a small pigeon hole principle instance.

Speaker: **Norbert Manthey** (TU Dresden)
Title: *Recent developments in parallel SAT solving*
Abstract:

Since 2008 most parallel SAT solvers follow the portfolio approach to run diverse solvers on the same formula. A formula is solved as soon as the fastest solver found a solution. This approach can be enhanced with sharing information, namely learned clauses, or the literals that have been used in the most recent conflict analysis. Orthogonally, different simplification methods can be applied within each solver, but to ensure soundness clause sharing needs to be restricted. Another development is to partition the search space by adding extra constraints to the formulas that are given to the different solvers. Again, clause sharing and simplifications can be implied - however, to obtain a sound procedure both sharing and simplification have to take care of the extra constraints. From a theoretical point of view the search space partitioning seems to scale better. In competitions, the most widely used approach is still portfolio based. The talk will give an overview about the two approaches, how to share clauses and the relation to formula simplifications during search.

Speaker: **Joao Marques Silva** (IST/INESC-ID)
Title: *Problem solving with SAT oracles*
Abstract:

The success of SAT solving is underscored by the widespread use of SAT solvers in practical applications. Whereas many practical applications can be cast as decision problems, for which a single query to a SAT oracle suffices, for many other applications, SAT oracles are called multiple times. For example, this is the case when solving decision problems with abstraction refinement, e.g. bit-vector formulas in SMT. This is also the case when solving decision problems in higher levels of the polynomial hierarchy, e.g. QBF. Moreover, many computational problems are naturally formulated as function (or search) problems, and can be solved with a number of queries to a SAT oracle. Concrete examples include computing a minimal unsatisfiable subformula, computing a maximum satisfied subformula, computing a prime implicate, computing a minimal model, among many other problems. This talk overviews problem-solving based on multiple queries to a SAT oracle, focusing on approaches for solving function problems. The talk presents a number of function problems defined on Boolean formulas and shows how most of these problems can be reduced to the problem of computing a minimal set subject to a monotone predicate (MSMP). The talk also describes a number of algorithms for the MSMP problem, highlighting the worst-case number of SAT oracle queries. Finally, the talk outlines a number of research topics in the area of problem solving with SAT oracles.

Speaker: **Nina Narodytska** (University of Toronto)
Title: *Reactive synthesis via QBF solving*
Abstract:

Reactive Software Synthesis is a growing research field that holds the promise of automating programmers' labour in domains such as device driver development, robotics, traffic control, etc. A reactive synthesis problem is formulated as a game between the system and its environment, where in order to win the system must correctly respond to any possible combination of inputs from the environment, including erroneous and malicious inputs.

In this talk I will present our ongoing research on developing a domain-specific Quantified Boolean Formula (QBF) solver for use in reactive synthesis. I will point out that reactive synthesis problems can be naturally formalised using QBF. However, modern QBF solvers can not solve even toy reactive synthesis problems. In this work, we aim to build a new QBF solver that can be used in synthesis of real-world

software. The specific class of reactive software that motivates our research are Operating System device drivers.

We build upon recently proposed RAReQS QBF solver that uses abstraction in order to transform a QBF problem into a sequence of SAT problems. We identify two main issues that lead to poor performance of existing QBF solvers on reactive synthesis problems, namely limited look-ahead and symmetric strategies subtrees. Both issues manifest in unnecessarily large search trees. To address these issues we propose to use stronger inference techniques at each step and keep a compact representation of strategy, like a DAG rather than tree.

This is joint work in progress with Fahiem Bacchus, Alexander Legg, Leonid Ryzhyk, Michael Stumm, and Adam Walker.

Speaker: **Jakob Nordström** (KTH Royal Institute of Technology)
Title: *Mini-tutorial on weak proof systems and connections to SAT solving*
Abstract:

This talk will focus on some comparatively weak proof systems that are being and/or could be used as a basis for SAT solvers. We will talk about resolution, polynomial calculus, and cutting planes (related to CDCL, Gröbner basis computations, and pseudo-Boolean solvers, respectively) and some proof complexity measures that have been studied for these systems, and briefly touch on if and how these complexity measures could be relevant for SAT solver performance. We will also discuss work on understanding how efficiently CDCL solvers can explore the search space of resolution proofs.

Speaker: **Albert Oliveras** (Universitat Politecnica de Catalunya)
Title: *Survey of satisfiability modulo theories (SMT)*
Abstract:

During the last fifteen years, mostly because of their efficiency and ease of use, SAT solvers have become the tool of choice for many diverse applications. However, propositional logic is very often not the most appropriate language to express the constraints that show up in real-world problems (e.g. arithmetic constraints). Satisfiability Modulo Theories (SMT) tries to address this problem by expressing the problems in a richer non-propositional logic, considering satisfiability of (usually ground) first-order formulas with respect to a background theory. In this talk, after a brief introduction to SMT, we will focus on DPLL(T), the algorithm that underlies all current state-of-the-art SMT solvers. We will describe its basic setting, some trivial optimizations and some further extensions that allow one to solve a variety of problems.

Speaker: **Ashish Sabharwal** (IBM Watson Research Center)
Title: *Resolution and parallelizability: Barriers to the efficient parallelization of SAT solvers*
Abstract:

Recent attempts to create versions of Satisability (SAT) solvers that exploit parallel hardware and information sharing have met with limited success. In fact, the most successful parallel solvers in recent competitions were based on portfolio approaches with little to no exchange of information between processors. This experience contradicts the apparent parallelizability of exploring a combinatorial search space. We present evidence that this discrepancy can be explained by studying SAT solvers through a proof complexity lens, as resolution refutation engines. Starting with the observation that a recently studied measure of resolution proofs, namely depth, provides a (weak) upper bound to the best possible speedup achievable by such solvers, we empirically show the existence of bottlenecks to parallelizability that resolution proofs typically generated by SAT solvers exhibit. Further, we propose a new measure of parallelizability based on the best-case makespan of an ofine resource constrained scheduling problem. This measure explicitly accounts for a bounded number of parallel processors and appears to empirically correlate with parallel speedups observed in practice. Our ndings suggest that efcient parallelization of SATsolvers is not simply a matter of designing the right clause sharing heuristics; even in the best case, it can be – and indeed is – hindered by the structure of the resolution proofs current SAT solvers typically produce.

Speaker: **Rahul Santhanam** (University of Edinburgh)
Title: *Beating brute force search for QBF satisfiability*
Abstract:

We give the first algorithms beating brute force search for general quantified Boolean formulas. We show that quantified CNFs over n variables of size $poly(n)$ with q alternations can be solved in time $2^{n-n^{1/(q+1)}}$, and also with a different algorithm in time $2^{n-\Omega(q)}$. This gives non-trivial savings for all values of q except in the regime $q \in [\theta(log(n)/loglog(n)), \theta(log(n))]$. We also show that improvements on our algorithms would lead to new lower bounds for Boolean formulas.

This is joint work with Ryan Williams.

Speaker: **Martina Seidl** (Johannes Kepler University)
Title: *Recent trends in QBF solving*
Abstract:

Quantified Boolean formulas (QBF) provide a powerful framework for efficiently encoding application problems located in PSPACE including problems from formal verification, artificial intelligence, etc. The language of QBF extends propositional logic by universal and existential quantifiers over the propositional variables. The interest in QBF is raised by the vision of using QBF solvers as effectively as SAT solvers.

In this talk, after a short overview of the history of QBF research, state-of-the-art techniques are presented which will presumably contribute to making QBF technology ready for practical applications. Amongst others, this includes preprocessing and certificate extraction. We outline how those techniques integrate into the solving workflow and how they affect the solving performance.

Speaker: **Laurent Simon** (University of Paris Sud)
Title: *Understanding the power of glue clauses*
Abstract:

In this talk, we will focus on the introduction of the concept of glue clauses and Literal Bock Distance in CDCL solvers. This measure for the quality of learnt clauses was introduced in 2009 and is now used in most of CDCL solvers. However, despite its interest, this measure is not fully understood. We will present the concept of glue clauses, as it was stated five years ago, and develop new insights in what may explain its performances, for instance by trying to find correlations between blocks as stated in the LBD measure and communities.

Speaker: **Carsten Sinz** (Karlsruhe Institute of Technology)
Title: *Abstraction and multi-encodings in SAT*
Abstract:

Abstraction is an extremely useful technique to handle complex systems, and has been applied in many areas such as model checking and SMT solving. Even though many applications of SAT solvers use abstraction (or abstraction refinement) techniques on a higher problem-level, they less frequently employ it on the clause level, and if so, mostly in connection with a particular problem at hand (e.g., bounded model checking).

In this talk I will summarize existing abstraction techniques in SAT solving and propose new directions on how they could be used in a general, problem-independent setting. I will also present ideas how to combine abstraction with multiple simultaneous problem encodings.

Speaker: **Stefan Szeider** (Vienna University of Technology)
Title: *Survey of parameterized complexity and SAT*
Abstract:

I will provide a gentle introduction to the field of parameterized complexity, with a special focus on satisfiability. The aim is to explain the different types of questions one can ask and to briefly summarize the known results without going much into the technical details.

Speaker: **Moshe Vardi** (Rice University)
Title: *Phase transitions and computational complexity*
Abstract:

In the past 20 years there has been extensive research exploring the relationship between the statistical behavior of random combinatorial problems and their computational complexity. This relationship also played a key role in recent attempts to resolve the status of P vs NP. A specific focus of this research has been on random 3-SAT, a paradigmatic NP-complete problem. An underlying assumption of this line of research has been that there is a fundamental connection between the satisfiability phase transition observed in random 3-SAT and and the average-case complexity of random 3-SAT. The common belief, for example, has been that hard random 3-SAT instances are located at the phase-transition region.

In this talk we offer a skeptical look at this line of research and question its basic assumptions. We demonstrate that there is no evidence for a fundamental connection between statistical behavior and intrinsic computational complexity. Rather, there seems to be a connection between statistical behavior and the performance of specific algorithms. This connection, however, is quite complicated and defies a unifying explanation.

Speaker: **Sean Weaver** (US Department of Defense)
Title: *Satisfiability-based set membership filters*
Abstract:

This presentation introduces a novel application of Satisfiability (SAT) to the set membership problem with specific focus on efficiently testing whether large sets contain a given element. Such tests are pervasive in modern computing (ex. database lookups, virus scanning, network routing, malicious website detection, ...) and can be greatly enhanced via the use of filters, probabilistic algorithms that can quickly decide whether or not a given element is in a given set. This talk will introduce SAT filters (i.e., filters based on SAT) and their use in the set membership problem. This presentation will show both the theoretical advantages of SAT filters and experimental results demonstrating that this technique yields significant performance improvements over previous techniques (such as the Bloom filter). Specifically, a SAT filter is a filter construction that is simple yet efficient in terms of both query time and filter size; i.e., SAT filters asymptotically achieve the information-theoretic limit while providing fast querying. As well, this is the first application that makes use of the random k-SAT phase transition results and may drive research into efficient solvers for this and similar applications.