

Research Summary (2018-2020)

Kuldeep S. Meel

Sung Kah Kay Assistant Professor, School of Computing

ENABLING BEYOND NP REVOLUTION

My research program's long-term vision is to advance automated reasoning techniques to enable computing to deal with increasingly uncertain real-world environments. From the core technical perspective, my research program is motivated by the unprecedented advances in the combinatorial solving techniques over the past three decades, often called as *NP revolution*, owing to the focus on the fundamental problem in Nondeterministic Polynomial (NP) time complexity: satisfiability (SAT). The NP revolution offers us the opportunity to develop scalable techniques for problems that lie *Beyond NP*. In particular, I focus on four fundamental problems that lie beyond NP: constrained counting, constrained sampling, functional synthesis, and maximum satisfiability. My research program is focused on enabling *Beyond NP Revolution* via making progress on the following critical enabling ingredients:

A. Beyond Worst-Case Algorithmic Design While theoreticians have studied the aforementioned problems for over thirty years, such studies employed the narrow lens of SAT oracles, which did not broadly allow a more nuanced characterization of the SAT queries' complexity. In contrast, we focus on designing algorithms with SAT queries that can be efficiently solved by the underlying SAT solvers. In the context of approximate counting, we have developed a hashing-based framework, ApproxMC, that relies on sparse XORs to circumvent the practical hardness of queries based on dense XORs [MSV19a, MSV19b, ABM20, MA20, BM20]. The LICS-20 paper [MA20] (co-authored with S. Akshay) resolved the longstanding open problem of designing a sparse XOR-based algorithmic framework that can achieve runtime improvement without losing theoretical guarantees. As a high degree of symmetry in the underlying constraints amplifies the practical hardness, we rely on domain-level symmetry breaking for performance improvement [WUA⁺20]. In the context of sampling, our framework, Waps, employs the progress in knowledge compilation to achieve scalability without losing theoretical guarantees [SGRM18, GSRM19]. In follow-up work, we improved the underlying algorithmic framework for knowledge compilation via probabilistic caching enabled via universal hashing [SRSM19]. Building on the success of constrained sampling and machine learning techniques, we proposed a novel algorithmic framework, called Manthan, for functional synthesis that achieved significant breakthrough in the number of instances solved [GRM20a].

B. Applications A key ingredient in the NP revolution was focus on the instances arising from the real-world, allowing the design of techniques tuned to practical instances. In a similar vein, we focus on discovering new applications for the aforementioned problems and the design of efficient modeling schemes. In an interdisciplinary collaboration, we reduced the problem of reliability of power transmission grids to constrained counting [PDOMV19]. Furthermore, we reduce the quantitative verification of neural networks [NSM⁺19, BSS⁺19], and quantification of information flow [BEH⁺18] to constrained counting. Building on our algorithmic progress in constrained sampling, we design efficient technique for higher t -wise coverage in the context of highly configurable software systems [BLM20]. In regards to maximum satisfiability, we showed that the problem of decoding in group testing reduces to maximum satisfiability [CGSM20]. In another line of work, we design efficient provable techniques for learning interpretable classifiers via reduction to MaxSAT [MM18, GM19, GMM20].

C. Query-Driven Design of Solvers SAT solvers' usage inside broader algorithmic frameworks gives rise to the rich structure of the SAT queries. We have taken a twofold approach: On one hand, our recent work on uncovering the structure of solution spaces through the lens of phase transition behavior [PJM19, GRM20b]. As a next step, we focus on engineering the underlying architecture of the SAT solvers. Given the crucial importance of CNF-XOR solving for counting and sampling, we have proposed a novel architecture, BIRD, relying on the tight integration of CDCL and Gauss-Jordan Elimination to achieve significant improvements for CNF-XOR solving [SM19, SGM20]. Motivated by applications in cryptography and collaboration with the Defense Service Organization (DSO), we developed a new architecture, Bosphorus, for CNF-ANF formulas [CSCM19], and further improved performance with phase saving [SM20]. Our most ambitious project, CrystalBall [SKM19], focuses on the data-driven synthesis of SAT Solvers by combining the advances in proof systems with supervised learning.

D. Testing Framework Similar to the pivotal role played by testing frameworks in the *NP revolution*, the realization of the Beyond NP revolution would need scalable testing and verification methodologies. As a first step, we have focused on testing of constrained samplers. In contrast to conventional programs, one trace/sample is typically inadequate to prove the non-conformity of the underlying samplers. We are developing an algorithmic framework, Barbarik, that sits at the intersection of property testing and symbolic reasoning [CM19]. Barbarik [CM19, MPC20] can handle product distributions and is the first tester to require a constant number of samples, in contrast to prior work requiring exponentially many samples.

Open-Source Tools Our research has led to the release of 10 actively maintained [open source tools](#). Our SAT solver entry won 3rd place in the Main Track of the highly prestigious and competitive SAT competition 2020 (the first instance of a top-3 finish by an entry from Singapore) while our model counting entry won 1st place in two of three tracks.

External Funding As a PI, I have attracted SGD 3.76 million of external funding (my share: SGD 3.25 million): SGD 2.6M from NRF Fellowship for AI (sole PI), SGD 180K from Defense Service Organization (as Sole PI), SGD 499K (co-PI: R. Yap) from AI Singapore, SGD 500K from NUS Reimagine Research (co-PI: B. Sikdar). Besides, I am one of the six co-PIs on the interdisciplinary [INTPREP](#) program.

REFERENCES

- [ABM20] Durgesh Agrawal, Bhavishya, and Kuldeep S. Meel. On the sparsity of xors in approximate model counting. In *Proceedings of the International Conference on Theory and Applications of Satisfiability Testing (SAT)*, 7 2020.
- [BEH⁺18] Fabrizio Biondi, Michael Enescu, Annelie Heuser, Axel Legay, Kuldeep S. Meel, and Jean Quilbeuf. Scalable approximation of quantitative information flow in programs. In *Proceedings of International Conference on Verification, Model Checking, and Abstract Interpretation (VMCAI)*, 1 2018.
- [BLM20] Eduard Baranov, Axel Legay, and Kuldeep S. Meel. Baital: An adaptive weighted sampling approach for improved t-wise coverage. In *Proceedings of ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE)*, 11 2020.
- [BM20] Jaroslav Bendik and Kuldeep S. Meel. Approximate counting of minimal unsatisfiable subsets. In *Proceedings of International Conference on Computer-Aided Verification (CAV)*, 7 2020.
- [BSS⁺19] Teodora Baluta, Shiqi Shen, Shweta Shine, Kuldeep S. Meel, and Prateek Saxena. Quantitative verification of neural networks and its security applications. In *Proceedings of ACM Conference on Computer and Communications Security (CCS)*, 11 2019.
- [CGSM20] Lorenzo Ciampiconi, Bishwamitra Ghosh, Jonathan Scarlett, and Kuldeep S. Meel. A MaxSAT-based framework for group testing. In *Proceedings of AAAI Conference on Artificial Intelligence (AAAI)*, 2 2020.
- [CM19] Sourav Chakraborty and Kuldeep S. Meel. On testing of uniform samplers. In *Proceedings of AAAI Conference on Artificial Intelligence (AAAI)*, 1 2019.
- [CMV13] Supratik Chakraborty, Kuldeep S. Meel, and Moshe Y. Vardi. A scalable approximate model counter. In *Proceedings of International Conference on Constraint Programming (CP)*, pages 200–216, 9 2013.
- [CSCM19] Davin Choo, Mate Soos, Kian Ming A. Chai, and Kuldeep S. Meel. Bosphorus: Bridging anf and cnf solvers. In *Proceedings of Design, Automation, and Test in Europe (DATE)*, 3 2019.
- [GM19] Bishwamitra Ghosh and Kuldeep S. Meel. Imli: An incremental framework for maxsat-based learning of interpretable classification rules. In *Proceedings of AAAI/ACM Conference on AI, Ethics, and Society (AIES)*, 1 2019.
- [GMM20] Bishwamitra Ghosh, Dmitry Malioutov, and Kuldeep S. Meel. Classification rules in relaxed logical form. In *Proceedings of European Conference on Artificial Intelligence (ECAI)*, 6 2020.
- [GRM20a] Priyanka Golia, Subhajit Roy, and Kuldeep S. Meel. Manthan: A data-driven approach for boolean function synthesis. In *Proceedings of International Conference on Computer-Aided Verification (CAV)*, 7 2020.
- [GRM20b] Rahul Gupta, Subhajit Roy, and Kuldeep S. Meel. Phase transition behavior in knowledge compilation. In *Proceedings of International Conference on Constraint Programming (CP)*, 9 2020.
- [GSRM19] Rahul Gupta, Shubham Sharma, Subhajit Roy, and Kuldeep S. Meel. Waps: Weighted and projected sampling. In *Proceedings of Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, 4 2019.

- [MA20] Kuldeep S. Meel and S. Akshay. Sparse hashing for scalable approximate model counting: Theory and practice. In *Proceedings of Logic in Computer science (LICS)*, 7 2020.
- [MM18] Dmitry Malioutov and Kuldeep S. Meel. Mlic: A maxsat-based framework for learning interpretable classification rules. In *Proceedings of International Conference on Constraint Programming (CP)*, 08 2018.
- [MPC20] Kuldeep S. Meel, Yash Pote, and Sourav Chakraborty. On testing of samplers. In *Proceedings of Advances in Neural Information Processing Systems(NeurIPS)*, 12 2020.
- [MSV18] Kuldeep S. Meel, Aditya A. Shrotri, and Moshe Y. Vardi. Not all fprass are equal: Demystifying fprass for dnf-counting. In *Proceedings of International Conference on Constraint Programming (CP)*, 08 2018. Invited to Constraints journal.
- [MSV19a] Kuldeep S. Meel, Aditya A. Shrotri, and Moshe Y. Vardi. Not all fprass are equal: Demystifying fprass for dnf-counting. *Constraints*, 08 2019.
- [MSV19b] Kuldeep S. Meel, Aditya A. Shrotri, and Moshe Y. Vardi. Not all fprass are equal: Demystifying fprass for dnf-counting (extended abstract). In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI)*, 08 2019. Sister Conference Best Papers Track.
- [NSM⁺19] Nina Narodytska, Aditya Shrotri, Kuldeep S. Meel, Alexey Ignatiev, and Joao Marques Silva. Assessing heuristic machine learning explanations with model counting. In *Proceedings of the International Conference on Theory and Applications of Satisfiability Testing (SAT)*, 7 2019.
- [PDOMV19] Roger Paredes, Leonardo Duenas-Osorio, Kuldeep S. Meel, and Moshe Y. Vardi. Principled network reliability approximation: A counting-based approach. *Journal of Reliability Engineering and System Safety(RESS)*, 11 2019.
- [PJM19] Yash Pote, Saurabh Joshi, and Kuldeep S. Meel. Phase transition behavior of cardinality and xor constraints. In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI)*, 8 2019.
- [SGM20] Mate Soos, Stephan Gocht, and Kuldeep S. Meel. Tinted, detached, and lazy cnf-xor solving and its applications to counting and sampling. In *Proceedings of International Conference on Computer-Aided Verification (CAV)*, 7 2020.
- [SGRM18] Shubham Sharma, Rahul Gupta, Subhajit Roy, and Kuldeep S. Meel. Knowledge compilation meets uniform sampling. In *Proceedings of International Conference on Logic for Programming Artificial Intelligence and Reasoning (LPAR)*, 11 2018.
- [SKM19] Mate Soos, Raghav Kulkarni, and Kuldeep S. Meel. Crystalball: Gazing in the black box of sat solving. In *Proceedings of the International Conference on Theory and Applications of Satisfiability Testing (SAT)*, 7 2019.
- [SM19] Mate Soos and Kuldeep S. Meel. Bird: Engineering an efficient cnf-xor sat solver and its applications to approximate model counting. In *Proceedings of AAAI Conference on Artificial Intelligence (AAAI)*, 1 2019.
- [SM20] Arijit Shaw and Kuldeep S. Meel. Designing new phase selection heuristics. In *Proceedings of the International Conference on Theory and Applications of Satisfiability Testing (SAT)*, 7 2020.
- [SRSM19] Shubham Sharma, Subhajit Roy, Mate Soos, and Kuldeep S. Meel. Ganak: A scalable probabilistic exact model counter. In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI)*, 8 2019.
- [WUA⁺20] Wenxi Wang, Muhammad Usman, Alyas Almaawi, Kaiyuan Wang, Kuldeep S. Meel, and Sarfraz Khurshid. A study of symmetry breaking predicates and model counting. In *Proceedings of Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, 4 2020.