Title: Some results on uniform generation and model counting

Given a propositional logic formula F in conjunctive normal form and a source of random bits, we consider two related problems: (i) generating models of F uniformly at random, and (ii) counting the number of models. If we could list all models of F explicitly, both problems are trivial. In practical applications, however, it is infeasible to list all models of F. Therefore, the above problems need to be solved without generating all models in the first place. Both problems have been of significant interest in both the theoretical and applied computer science communities. For instance, propositional model-counting is the canonical #P-complete problem, and "almost uniform" generation is known to be polynomially inter-reducible to approximate model counting. As far as applications go, solutions to the above problems find use in verification, testing and probabilistic inference and analysis. Earlier approaches to these problems belong to one of two distinct categories: (i) theoretically motivated approaches that provide strong guarantees but do not lend themselves to efficient implementations in practice, or (ii) practically motivated approaches that scale to large problem sizes but provide very weak guarantees. In this talk, we describe our ongoing work on bridging these extremes. Specifically, we show that universal hashing functions can be used in a key way to achieve strong theoretical guarantees (similar to the best known guarantees from theoretically motivated approaches) and also scale to practical problem sizes (comparable to those achieved by practically motivated approaches). Our approach poses some unanswered questions on the relation between relaxed notions of uniform model generation and approximate model counting.

This is joint work with Kuldeep Singh Meel and Moshe Vardi