Four Subareas of the Theory of Constraints, and Their Links

Albert Atserias

Universitat Politècnica de Catalunya, Barcelona, Spain

Let $V = \{x_1, \ldots, x_n\}$ be a set of variables that range over a set of values $D = \{d_1, \ldots, d_q\}$. A constraint is an expression of the type $R(x_{i_1}, \ldots, x_{i_r})$, where $R \subseteq D^r$ is a relation on the domain set D and $x_{i_1}, \ldots, x_{i_r}$ are variables in V. The space of assignments, or configurations, is the set of all mappings $\sigma : V \to D$. We say that σ satisfies the constraint $R(x_{i_1}, \ldots, x_{i_r})$ if $(\sigma(x_{i_1}), \ldots, \sigma(x_{i_r})) \in R$. Otherwise we say that it falsifies it. On a given *system of constraints* we face a number of important computational problems. Here is a small sample:

1. Is it satisfiable? That is, is there an assignment that satisfies all constraints? If it is satisfiable, can we find a satisfying assignment? If it is unsatisfiable, does it have a small and efficiently checkable proof of unsatisfiability? Can we find it?
2. If it is not satisfiable, what is the maximum number of constraints that can be satisfied simultaneously by some assignment? Knowing that we can satisfy a $1 - \epsilon$ fraction of the constraints simultaneously for some small $\epsilon > 0$, can we find an assignment that satisfies more than, say, a $1 - \sqrt{\epsilon}$ fraction?
3. How many satisfying assignment does the system have? If we can't count it exactly, can we approximate the number of satisfying assignments up to a constant approximation factor? Can we sample a satisfying assignment uniformly or approximately uniformly at random? More generally, if we write $H(\sigma)$ for the number of constraints that are falsified by σ, can we sample an assignment σ with probability proportional to $e^{-\beta H(\sigma)}$ where β is a given *inverse temperature* parameter? Can we compute, exactly or approximately, the so-called partition function of the system, defined as $Z(\beta) = \sum_\sigma e^{-\beta H(\sigma)}$?
4. Knowing that the system has been generated randomly by choosing each R uniformly at random from a fixed set of relations Γ and by choosing each $(x_1, \ldots, x_r)$ uniformly at random in V^r, can we analyze and exploit the typical structure of the assignment space and the constraint system to solve any of the problems above?

These are fundamental problems that can be roughly classified into the following four categories: logic and proof complexity, optimization and approximation, counting and sampling, and analysis of randomly generated instances. Perhaps surprisingly, these four areas of the theory of constraints have been approached through rather different techniques by groups of researchers with small pairwise intersections. However, recent work has shown that these areas might be more related than this state of affairs seems to indicate. We overview these connections with emphasis on the open problems that have the potential of making the connections even tighter.

R. Královič and D. Niwiński (Eds.): MFCS 2009, LNCS 5734, p. 1, 2009.