

Abstract

An instance of the Constraint Satisfaction Problem (CSP) is given by a family of constraints on overlapping sets of variables, and the goal is to assign values from a fixed domain to the variables so that all constraints are satisfied. In the optimization version, the goal is to maximize the number of satisfied constraints. An approximation algorithm for CSP is called robust if it outputs an assignment satisfying a $(1 - g(\varepsilon))$ -fraction of constraints on any $(1 - \varepsilon)$ -satisfiable instance, where the loss function g is such that $g(\varepsilon) \rightarrow 0$ as $\varepsilon \rightarrow 0$. We study how the robust approximability of CSPs depends on the set of constraint relations allowed in instances, the so-called constraint language. All constraint languages admitting a robust polynomial-time algorithm (with some g) have been characterised by Barto and Kozik, with the general bound on the loss g being doubly exponential, specifically $g(\varepsilon) = O((\log \log(1/\varepsilon))/\log(1/\varepsilon))$. It is natural to ask when a better loss can be achieved: in particular, polynomial loss $g(\varepsilon) = O(\varepsilon^{1/k})$ for some constant k . In this paper, we consider CSPs with a constraint language having a near-unanimity polymorphism. We give two randomized robust algorithms with polynomial loss for such CSPs: one works for any near-unanimity polymorphism and the parameter k in the loss depends on the size of the domain and the arity of the relations in Γ , while the other works for a special ternary near-unanimity operation called dual discriminator with $k = 2$ for any domain size. In the latter case, the CSP is a common generalisation of UNIQUE GAMES with a fixed domain and 2-SAT. In the former case, we use the algebraic approach to the CSP. Both cases use the standard semidefinite programming relaxation for CSP.