


# DISTRIBUTED CONSTRAINT OPTIMIZATION: PRIVACY GUARANTEES AND STOCHASTIC UNCERTAINTY

Thèse n. 5197 2011  
présentée le 29 septembre 2011  
à la Faculté Informatique et Communications  
Laboratoire d'Intelligence Artificielle  
programme doctoral en Informatique et Communications  
École Polytechnique Fédérale de Lausanne

pour l'obtention du grade de Docteur ès Sciences  
par

Thomas Léauté



acceptée sur proposition du jury :

Prof. Karl Aberer, président du jury  
Prof. Boi Faltings, directeur de thèse  
Prof. Jean-Pierre Hubaux, rapporteur  
Prof. Makoto Yokoo, rapporteur  
Dr. Juan A. Rodríguez-Aguilar, rapporteur

Lausanne, EPFL, 2011

# Abstract

*Distributed Constraint Satisfaction (DisCSP)* and *Distributed Constraint Optimization (DCOP)* are formal frameworks that can be used to model a variety of problems in which multiple decision-makers cooperate towards a common goal: from computing an equilibrium of a game, to vehicle routing problems, to combinatorial auctions. In this thesis, we independently address two important issues in such multi-agent problems: 1) how to provide strong guarantees on the protection of the privacy of the participants, and 2) how to anticipate future, uncontrollable events.

On the privacy front, our contributions depart from previous work in two ways. First, we consider not only *constraint privacy* (the agents' private costs) and *decision privacy* (keeping the complete solution secret), but also two other types of privacy that have been largely overlooked in the literature: *agent privacy*, which has to do with protecting the identities of the participants, and *topology privacy*, which covers information about the agents' co-dependencies. Second, while previous work focused mainly on quantitatively measuring and reducing privacy loss, our algorithms provide stronger, qualitative guarantees on what information will remain secret. Our experiments show that it is possible to provide such privacy guarantees, while still scaling to much larger problems than the previous state of the art.

When it comes to reasoning under uncertainty, we propose an extension to the DCOP framework, called *DCOP under Stochastic Uncertainty (StochDCOP)*, which includes uncontrollable, random variables with known probability distributions that model uncertain, future events. The problem becomes one of making "optimal" offline decisions, before the true values of the random variables can be observed. We consider three possible concepts of optimality: minimizing the *expected* cost, minimizing the *worst-case* cost, or maximizing the probability of a-posteriori optimality. We propose a new family of StochDCOP algorithms, exploring the tradeoffs between solution quality, computational and message complexity, and privacy. In particular, we show how discovering and reasoning about co-dependencies on common random variables can yield higher-quality solutions.

**Keywords:** Distributed Constraint Satisfaction (DisCSP), Distributed Constraint Optimization (DCOP), Privacy, Uncertainty, P-DPOP,  $\mathbb{E}$ [DPOP], FRODO