

Abstract

This book focuses on Distributed Constraint Optimization as an approach to coordinate the actions of a network of cooperative agents. Different scenarios pose very different problems if we intend to endow agents with coordinated behaviour due to availability of resources. Some resource-bounded problems require complete DCOP algorithms that can find the most effective use of resources to achieve an optimal coordination, while others need incomplete algorithms that sacrifice optimality in favour of low-cost suboptimal solutions.

This motivates the main goal of this work: the design of efficient DCOP algorithms to cope with different resource-bounded scenarios while ensuring that agents select the actions that allow their coordination. Quality assessment is likely to play an important role in this endeavour, because it is fundamental to weigh the cost of coordination against the quality of the solution reached (trade-off quality versus cost). Unfortunately, quality assessment for incomplete DCOP algorithms is a major challenge that requires a broader concept of guarantees that includes approximate quality guarantees.

This book addresses these two major issues: efficiency and approximate quality assessment. We make significant contributions to both issues. The main idea behind these contributions is to design algorithms, and frameworks, which exploit a DCOP structure, namely the structure of agents dependencies and of their rewards, to assess and bound high quality solutions.

Regarding optimality guarantees, we contribute with Action-GDL, a complete DCOP algorithm that generalises and improves the efficiency of existing dynamic programming DCOP algorithms, unifying them under the GDL framework. The key idea behind Action-GDL is to better exploit a DCOP structure by using a novel problem representation based on junction trees.

As to approximate quality assessment, we propose two general frameworks: Divide-and-Coordinate (DaC) and Region Optimality.

First, the DaC framework enables to trade-off quality versus cost from an agent perspective by defining a family of algorithms, the DaC family, which allows agents to bound the quality of their solutions at run time. The DaC framework defines a new approach to DCOP solving by: (i) splitting the problem into simpler sub-problems that are computationally tractable; and (ii) having sub-problems agree on a solution. We define DaCSA and EU-DaC, two DaC-compliant, incomplete DCOP algorithms. Both algorithms can return anytime solutions with quality guarantees.

Secondly, this work introduces the region optimal framework, a general framework that

generalises local optimality frameworks introduced in the literature such as k-size or t-size. Region optimality allows to obtain quality guarantees for arbitrary criteria and depending on the available information about a DCOP. Moreover, we show that there are criteria that outperform state-of-the-art local criteria such as (k-)size or (t-)distance by introducing the novel size-bounded distance criterion. This contribution finishes by proposing \mathcal{C} -DALO, an asynchronous region optimal DCOP algorithm to search for region optimal solutions in any region characterised by any arbitrary criteria.

Finally, we prove that region optimality is a valuable tool for bounding the quality of the solutions achieved by the Max-Sum algorithm on convergence. These results shed light on the relationship between the Max-Sum performance and the structure of the problem. Moreover, they help identify new classes of graph structures for which Max-Sum is guaranteed to converge to high-quality solutions.