Expert Systems with Applications 42 (2015) 7797–7811

Contents lists available at ScienceDirect

Expert Systems with Applications

journal homepage: www.elsevier.com/locate/eswa

Parallel multi-objective multi-robot coalition formation

Manoj Agarwal [a], Nitin Agrawal [b], Shikhar Sharma [b], Lovekesh Vig [c,*], Naveen Kumar [a]

[a] Department of Computer Science, University of Delhi, Delhi 110007, India
[b] Cluster Innovation Centre, University of Delhi, Delhi 110007, India
[c] School of Computational and Integrative Sciences, Jawaharlal Nehru University, New Delhi 110067, India

ARTICLE INFO

Article history:
Available online 8 June 2015

Keywords:
CUDA
Multi-robot
Additive
Non-additive
Task allocation
Evolutionary algorithms
Multi-objective optimization
Coalition formation

ABSTRACT

In the quest for greater autonomy, there is an increasing need for solutions that would enable a large set of robots to coalesce and perform complicated multi-robot tasks. This problem, also known as the multi-robot coalition formation problem has been traditionally approached as a single objective optimization problem. However, robots in the real world have to optimize multiple conflicting criteria such as battery life, number of completed tasks, and distance traveled. Researchers have only recently addressed the robot coalition formation problem as a multi-objective optimization problem, however the proposed solutions have computational bottlenecks that make them unsuitable for real time robotic applications. In this paper we address the issue of scalability by proposing parallelized algorithms in the CUDA programming framework. NSGA-II and PAES algorithm have been parallelized due to their suitability to the coalition formation domain as outlined in our previous work. The parallelized versions of these algorithms have been applied to both the additive and non-additive coalition formation environments. Simulations have been performed in the player/stage environment to validate the applicability of our approach to real robot situations. Results establish that the multi-point PAES parallel variant yields significant performance gains in terms of running time and solution quality when the problem is scaled to deal with large inputs. This suggests that the algorithm may be viable for real time robotic applications. Experiments demonstrate significant speedup when the proposed parallel algorithms were compared with the serial solutions proposed earlier.

1. Introduction

Most real world problems have multiple conflicting objectives, i.e. improvement with respect to one objective leads to degradation with respect to one or more objectives. In such problems, no single solution can serve as the optimal solution with respect to all the objectives and therefore most multi-objective optimization techniques result in a set of trade-off solutions that are mutually non-dominant.[1] Such a set of mutually non-dominated solutions is known as the Pareto optimal set. Although several aggregation techniques to convert multi-objective optimization problems into single objective optimization problems have been proposed in the literature, most work (Deb, Mohan, & Mishra, 2003, 2002; Corne, Jerram, Knowles, Oates, & Martin, 2001, 2000; Knowles & Corne, 1999; Kim, Hiroyasu, Miki, & Watanabe, 2004; Zitzler & Thiele, 1999) in the multi-objective optimization domain has focused on approximating the Pareto optimal set of solutions. Metaheuristic techniques, such as evolutionary approaches, have established their supremacy for such problems due to their problem independent search capabilities. Though these techniques do not guarantee the Pareto optimal set of solutions as the end result, they produce solutions that are very good approximations of the Pareto optimal set. However, approximating the Pareto optimal set of solutions from a large and complex search space can be computationally expensive depending on the problem at hand. As a result, parallelization of these algorithms becomes important for effective deployment.

The multi-robot coalition formation problem deals with the formation of multi-robot teams that can be assigned to a set of given tasks for their execution. This problem is a member of the NP-hard class of combinatorial optimization problems in which the goal is to find the optimal mapping of robots to task specific teams. Task allocation is a well studied problem in the multi-agent domain and is gaining importance in the robotics community as the

* Corresponding author.
E-mail addresses: agar.manoj@gmail.com (M. Agarwal), nitin.cic@gmail.com (N. Agrawal), shikhar.cic23@gmail.com (S. Sharma), lovekesh@jnu.ac.in (L. Vig), nk.cs.du@gmail.com (N. Kumar).

[1] The concept of dominance between two solutions, $x^{(1)}$ and $x^{(2)}$, is as follows: A solution $x^{(1)}$ is said to dominate another solution $x^{(2)}$, if both of the following are true:
1. $x^{(1)}$ is no worse than $x^{(2)}$ in all the objectives.
2. $x^{(1)}$ is strictly better than $x^{(2)}$ in at least one of the objectives.Solutions are called non-dominated if they do not dominate each other.

http://dx.doi.org/10.1016/j.eswa.2015.05.032
0957-4174/

complexity of robotic tasks increases. However, solutions from the multi-agent systems domain cannot be directly applied to the multi-robot systems (Vig & Adams, 2005). Gerkey and Mataric (2004) have categorized the domain of multi-robot task allocation along three mutually perpendicular axes: Single-Robot tasks (SR) vs Multi-Robot (MR) tasks (coalition/team dependent task), Single-Task (ST) robots vs Multi-Task (MT) robots, and Instantaneous Assignment (IA) vs Time-extended Assignment (TA). In this paper, the ST–MR–IA problem category, also known as the multi-robot coalition formation problem has been modeled as a multi-objective optimization problem using the CUDA parallel programming framework.

Recently, substantial performance gain in solving computation intensive problems has been achieved using the Compute Unified Device Architecture (CUDA) platform. Using the single instruction multiple thread (SIMT) technique, CUDA enables parallelism by calling kernel functions from sequential code that launch threads executing the same kernel functions in parallel. Thus computations can be implemented in a highly parallelized way yielding significant performance enhancement over contemporary single threaded machines.

Popular algorithms like NSGA-II, SPEA2 and PAES have previously been adapted to the problem of multi objective robot coalition formation (Agarwal, Kumar, & Vig, 2014). However the algorithms were not scalable and the computational time requirements made them unsuitable for real world robotic applications. In this paper the authors attempt to address the issue of time complexity encountered in the previous paper by developing parallelized versions of the algorithms in the CUDA programming framework. The proposed approach is then compared with the original serial versions of the algorithms in order to ascertain speedup. NSGA-II and PAES were parallelized due to their suitability to the coalition formation domain as established in the previous work. Additionally, the algorithms were implemented for both the additive and non-additive coalition formation environments. Experiments establish that the multi-point PAES parallel variant yields significant performance gains in terms of running time and solution quality when the problem is scaled to deal with large inputs. Using the player-stage robotic simulation environment,the applicability of the proposed approach to real world robotic tasks has been demonstrated.

The organization of the remaining paper is as follows: Section 2 describes the related literature, Section 3 briefly summarizes the important evolutionary multi-objective optimization concepts, Section 4 formally describes the problem statement and the proposed solution strategies, Section 5 gives experimental settings and the related results, and finally Section 6 summarizes conclusions and scope of future work.

2. Related work

Robot coalition formation is a well studied problem in the multi-robot domain and researchers have developed numerous heuristic based (Service & Adams, 2011; Vig & Adams, 2006b) and market based (Vig & Adams, 2006a; Zlot & Stentz, 2006) solutions. Most existing approaches to the coalition formation problem (Dias, 2004; Gerkey & Mataric, 2002; Jones et al., 2006; Li, Xu, Yang, Chen, & Li, 2009; Liu & Chen, 2006; Shiroma & Campos, 2009; Shehory & Kraus, 1998; Vig & Adams, 2006a) do not address the problem of task execution i.e. once the coalitions have been formed how do the robots actually perform the joint team tasks. Tang and Parker's ASyMTRe (Tang & Parker, 2005) system was the first system that decomposed a multi-robot task into schemas that are executable by individual robots. The algorithm is capable of generating multiple decompositions for a particular multi-robot task i.e. it succeeded in dividing the task in many different ways to be performed by different combinations of robots with varying efficiencies. As ASyMTRe is a system for task decomposition and execution, the outcome of ASyMTRe may be used as input by a coalition formation algorithm. In this paper we incorporate the task decompositions provided by ASyMTRe into our task definitions as input to our coalition formation process. In this way we demonstrate how the entire process from task decomposition to allocation can be completely automated. ASyMTRe was later extended to work in distributed environments and to work with tightly coupled tasks (Zhang & Parker, 2010).

Liu and Chen (2006) provide a genetic approach for searching the coalition structure space to obtain the optimal coalition structure. Recently quantum-inspired approaches to coalition optimization have gained popularity. In particular Li et al. (2009) propose a quantum evolutionary approach for coalition formation, where a skill based quantum probability representation of chromosome coding strategy is designed to adapt to the multiple robot coalition formation problem. Zhang, Liu, Fu, and Wu (2009) propose a quantum inspired ant colony optimization (QACO) method to improve the ability to search and optimize coalition structures. Each ant is a quantum individual and probabilities of choosing a particular robot are utilized to optimize the QACO problem.

More recently, Sen and Adams (2014) developed a system with a library of coalition formation strategies that intelligently selects the appropriate strategy on the fly based on domain dependent features via the use of influence diagrams. The diagrams are generated via prominent features detected via principal components analysis. Service, Sen, and Adams (2014) developed a system to address the issue of task preemption in dynamic environments. They propose a simultaneous descending auction to prevent unnecessary task reassignments in the previously described RACHNA (Vig & Adams, 2007) system during task preemption. Sen and Adams (2013) present a simulated annealing inspired ANT colony optimization algorithm for spatially distributed real-world tasks in multi-robot systems. However, they deal with only one objective – the distance traveled. Zhang, Parker, and Kambhampati (2014) provide a mechanism for coalition coordination with tightly coupled robotic tasks. The algorithm utilizes the IQ-ASyMTRe framework outlined in Zhang and Parker (2010) to determine the sensor constraints and utilize a constraint graph to reason about whether these constraints can be maintained during the execution of the multi-robot tasks. Zhong et al. (2014) implement an auction based protocol for coalition formation that forces robots to look at both health levels and task requirements before making a bid. The algorithm is sub-optimal and is shown to be both complete and sound. While the above solutions are significant steps in the direction of imposing robot coalition formation algorithms, all of them model the multi-robot coalition formation problem as a single-objective problem, and coalitions are evaluated based upon a pre-determined utility function, whereas the coalition formation problem belongs to the multi-objective optimization domain.

The formulation of the multi-robot coalition problem as a multi objective optimization problem has been investigated recently (Agarwal et al., 2014). The coalition formation problem inherently lends itself to the multi-objective approach due to the demands of real world robotic tasks which require optimization of multiple conflicting objectives such as maximization of battery life and maximization of distance traveled by the robot. Agarwal et al. (2014) implemented variants of the well known PAES, NSGA-II and SPEA2 algorithms to arrive at the Pareto front of potential coalition structures. However, these methods fail to scale up to a large number of robots. In this paper we formulate the coalition formation problem as one with multiple conflicting objectives. Furthermore, we have utilized evolutionary approaches, accelerated via parallelization in GPU framework to arrive at fast, high quality coalition schemes.