BEN-GURION UNIVERSITY OF THE NEGEV

Communication Awareness in Multi-Agent Optimization

Author:
Ben Rachmut

Supervisor:
Professor Roie Zivan

Thesis submitted in partial fulfillment of the requirements for the degree of "Doctor of Philosophy"

Submitted to the Senate of Ben-Gurion University of the Negev
March 31, 2024

Beer-Sheva

אוניברסיטת בן גוריון בנגב

מודעות תקשורת באופטימיזציה של מערכות מרובות סוכנים

מאת:
בן רחמוט

בהנחיית:
פרופ' רועי זיוון

מחקר לשם מילוי חלקי של הדרישות לקבלת תואר

"דוקטור לפילוסופיה"

הוגש לסינאט אוניברסיטת בן גוריון בנגב

כ"א באדר ב
תשפד

31 במרץ
2024

באר שבע

BEN-GURION UNIVERSITY OF THE NEGEV

Communication Awareness in Multi-Agent Optimization

Author:
Ben Rachmut

Supervisor:
Professor Roie Zivan

Thesis submitted in partial fulfillment of the requirements for the degree of "Doctor of Philosophy"

Submitted to the Senate of Ben-Gurion University of the Negev

Approved by the advisor: ______________

Approved by the Dean of the Kreitman School of Advanced Graduate Studies: ______________

March 31, 2024

Beer-Sheva

Declaration of Authorship

This work was carried out under the supervision of
Prof. Roie Zivan
In the Department of Industrial Engineering and Management
Faculty of Engineering Sciences

I, Ben Rachmut, whose signature appears below, hereby declare that this thesis titled, "Communication Awareness in Multi-Agent Optimization" and the work presented in it are my own. I confirm that:

- I have written this Thesis by myself, except for the help and guidance offered by my Thesis Advisors.
- The scientific materials included in this Thesis are products of my own research, culled from the period during which I was a research student.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given.

Student's name: Ben Rachmut

Signature:

Date:

Acknowledgements

I would like to acknowledge everyone who has assisted me throughout my doctoral studies over the years.

Abstract

Most studies investigating multi-agent optimization (MAO) assume that messages arrive instantaneously and are never lost. Distributed constraint optimization problems (DCOPs) and general task allocation problems (GTAPs) are examples of models used to solve real-life MAO problems; however, these models do not consider the effects of imperfect communication. In distributed systems, in which communication is a key factor, agents rely on message delivery to achieve mutual goals. However, assuming seamless communication is not always practical because it may be subject to delays and failures. To address this gap in existing models, we introduce communication-aware multi-agent optimization (CA-MAO) models. These models are specifically designed to incorporate dynamic uncertainties in communication within distributed systems. Extending models to their communication-aware versions enhances their ability to represent real-world distributed settings where communication challenges are taken into consideration.

Despite running in asynchronous environments, many state-of-the-art MAO algorithms are designed synchronously. These algorithms rely on synchronous iterations in which each agent interacts with each neighbor in each cycle; the challenge arises from the uncertainty in message delivery. The reliance on message delivery creates a significant challenge for such synchronous algorithmic designs. Uncertainty in message delivery can cause delays or even deadlocks because agents can get stuck waiting for messages that never arrive. These issues significantly affect the efficiency and effectiveness of the algorithm.

In addition, some of these algorithms and mechanisms exploit common simplistic communication assumptions to achieve desirable results, such as monotonicity and convergence. To address the limitations of current algorithms that do not consider the effects of imperfect communication, we propose several communication-aware algorithms.

This study advances the field of MAO research in the following four ways:

1. First, we explore distributed local search algorithms for CA-DCOPs. We evaluate the performance of existing synchronous algorithms and develop their asynchronous counterparts. We also propose a novel asynchronous-anytime framework to enhance exploration by nonmonotonic asynchronous algorithms. The results demonstrate that imperfect communication surprisingly improves distributed local search algorithms, which increases exploration.

2. The second contribution focuses on distributed local search algorithms that guarantee solutions that cannot be improved by a group of k agents, known as k-opt solutions. We propose the latency-aware monotonic distributed local search (LAMDLS) and LAMDLS-2

algorithms, which are latency-aware search algorithms to solve DCOPs. These algorithms are monotonic and guarantee 1- and 2-opt convergence, respectively. In addition, they are resilient to message latency.

3. Our third contribution consists of investigating inference algorithms for CA-DCOPs, with a particular focus on the max-sum algorithm and its variants. We assess the performance of synchronous and asynchronous versions of the algorithm in various communication scenarios through empirical and theoretical analyses. The results indicate that the damped max-sum algorithm is remarkably robust when faced with imperfect communication.

4. Finally, the fourth contribution is the introduction of the CA-GTAPs. We focus on the Fisher market-clearing task allocation (FMC_TA) algorithm, which outperforms centralized and distributed optimization algorithms but is unsuited for imperfect communication scenarios. We propose an asynchronous FMC_TA algorithm that demonstrates resistance to imperfect communication without compromising solution quality. Our investigation also shows that, when communication is extremely poor, the distributed version of the algorithm performs better than the centralized version and produces consistent results.

Keywords: Multi-Agent Optimization, Communication-Awareness, Distributed Constraint Optimization, Task Allocation, Asynchronous Algorithms, Imperfect Communication

תקציר

רוב המחקרים בתחום אופטימיזציה מרובת סוכנים (MAO) מניחים שהודעות מגיעות מיידית ולא הולכות לאיבוד. בעיות אופטימיזציה מבוזרות של אילוצים (DCOPs) ובעיות הקצאת משימות כלליות (GTAPs) הן דוגמאות למודלים המשמשים לפתרון בעיות MAO מציאותיות. עם זאת, מודלים אלו לא מתחשבים בהשפעות של הפרעות בתקשורת. במערכות מבוזרות, בהן תקשורת היא גורם מפתח, סוכנים מסתמכים על העברת הודעות כדי להשיג מטרות משותפות. עם זאת, ההנחה של תקשורת מושלמת לא תמיד מציאותית, שכן היא עלולה להיות נתונה לעיכובים וכשלים. כדי להתמודד עם פער זה במודלים הקיימים, אנו מציגים מודלים של אופטימיזציה מרובת סוכנים עם מודעות תקשורת מודלים אלה תוכננו במיוחד כדי לשלב חוסר ודאות דינמית בתקשורת, בתוך מערכות מבוזרות. הרחבת המודלים מאפשרת שיפור של יכולתם לייצג בעיות מבוזרות בעולם האמיתי המתחשבות באתגרי תקשורת.

רבים מהאלגוריתמים עבור פתרון בעיותMAO תוכננו באופן סינכרוני, למרות שהם פועלים בסביבה אסינכרונית. האלגוריתמים הללו מסתמכים על איטרציות סינכרוניות, בהן סוכן מתקשר עם כל שכן בכל מחזור. הסתמכות על העברת הודעות יוצרת אתגר משמעותי עבור עיצוב סינכרוני. אי ודאות בהעברת הודעות עלולה לגרום לעיכובים ואף למבוי סתום, שכן סוכנים עלולים להיתקע בהמתנה להודעות שלא יגיעו לעולם. בעיות אלו יכולות להשפיע באופן משמעותי על היעילות והאפקטיביות של האלגוריתם.

מחקר זה מקדם את תחום חקר ה-MAO בארבע דרכים:

1. חקר אלגוריתמי חיפוש מקומיים מבוזרים עבור בעיות אופטימיזציה מבוזרות עם אילוצים: הערכנו את הביצועים של האלגוריתמים הסינכרוניים הקיימים ופיתחנו את גרסאותיהם האסינכרוניות. בנוסף, אנו מציעים מסגרת חדשה, אסינכרונית, לקבלת תכונת ה,anytime על מנת לאפשר חקירה במרחב הפתרונות (exploration) עבור אלגוריתמים אסינכרוניים לא מונוטוניים. התוצאות שלנו מראות שתקשורת לא מושלמת יכולה לשפר באופן מפתיע אלגוריתמי חיפוש מקומיים מבוזרים באמצעות הגברת החקירה של מרחב הפתרונות.
2. התרומה השנייה שלנו מתמקדת באלגוריתמי חיפוש מקומיים מבוזרים, המבטיחים פתרונות שאינם ניתנים לשיפור על ידי קבוצה של K סוכנים, הידועים כפתרונות K-OPT. אנו מציגים את ה (Latency-Aware Monotonic Distributed Local Search (LAMDLS ואת 2-LAMDLS אלגוריתמי חיפוש עבור בעיות אופטימיזציה מבוזרות עם אילוצים ותקשורת מודעת. אלגוריתמים אלו מבטיחים התכנסות של opt-1 ו opt-2 בהתאמה.

3. התרומה השלישית שלנו מתמקדת בחקר אלגוריתמי הסקה עם דגש על אלגוריתם ה Max sum והגרסאות השונות שלו, עבור בעיות אופטימיזציה מבוזרות עם אילוצים בסביבה עם בעיות תקשורת. באמצעות ניתוחים אמפיריים ותיאורטיים, בחנו את הביצועים של גרסאות סינכרוניות ואסינכרוניות של האלגוריתם בתרחישי תקשורת שונים. הממצאים שלנו מצביעים על כך שאלגוריתם הdamp Max sum מפגין עמידות יוצאת דופן.
4. לבסוף הצגנו את בעיות הקצאת משימות הכללית עם בעיות תקשורת. אנו התמקדנו באלגוריתם Fisher Market Clearing Task Allocation המפגין ביצועים טובים יותר מאלגוריתמי אופטימיזציה ריכוזיים ומבוזרים רבים אך אינו מתאים לתרחישי תקשורת לא מושלמת. לצורך התמודדות עם אתגר זה, אנו מציעים אלגוריתם אסינכרוני בשם FMC_ATA המפגין חסינות לבעיות תקשורת ללא התפשרות על איכות הפתרון. המחקר שלנו גם מראה שהגרסה המבוזרת של האלגוריתם עדיפה על הגרסה הריכוזית, כאשר התקשורת גרועה מאוד, והיא מפיקה תוצאות עקביות.

מילות מפתח: אופטימיזציה מרובת סוכנים, מודעות לתקשורת, אופטימיזציה של אילוצים מבוזרים, הקצאת משימות, אלגוריתמים אסינכרוניים, תקשורת לא מושלמת.

Contents

List of Figures

List of Tables

List of Abbreviations

CA	**C**ommunication **A**ware
MAO	**M**ulti-**A**gent **O**ptimization
CCG	**C**onstrained **C**ommunication **G**raph
DALO	**D**istributed **A**synchronous **L**ocal **O**ptimization
DCOP	**D**istributed **C**onstraint **O**ptimization **P**roblem
DSA	**D**istributed **S**tochastic **A**lgorithm
DSA-SDP	**D**istributed **S**tochastic **A**lgorithm- **S**lope **D**ependent **P**robability
FMC_ATA	**F**isher **M**arket **C**learing **A**synchronous **T**ask **A**llocation
FMC_TA	**F**isher **M**arket **C**learing **T**ask **A**llocation
GTAP	**G**eneral **T**ask **A**llocation **P**roblem
LAMDLS	**L**atency **A**ware **M**onotonic **D**istributed **L**ocal **S**earch
MAS	**M**ulti-**A**gent **S**ystems
MGM	**M**aximum **G**ain **M**essages

Chapter 1

Introduction

Recent advances in computation and communication have resulted in realistic distributed applications in which humans and technology interact and strive to optimize mutual goals (e.g., Internet of Things applications). Thus, the demand is growing for optimization methods to support decentralized decision-making in complex multi-agent systems (MASs), which find applications in various domains, including device scheduling in smart homes, target tracking in sensor networks, mission planning for unmanned autonomous vehicles, and coordination of rescue units in disaster scenarios, as noted in Refs. [10, 22, 32].

Understanding the complexities and challenges associated with MAS requires understanding the disaster management domain. In disasters, diverse rescue units such as medical staff, firefighters, and law enforcement must coordinate their actions to save as many victims as possible. Such a multi-agent task coordination problem is challenging due to the following characteristics:

- **Optimization of a global objective.** The various rescue units must work together toward a common goal (e.g., saving the maximum number of victims). An example is when police units must ensure the safe passage of ambulances carrying victims out of a disaster area.

- **Decentralized coordination.** A centralized entity that manages the coordination problem is often lacking. For example, hospital administrators want to manage their medical personnel, fire chiefs want to manage their firefighters, and police chiefs want to manage their police units. If a larger incident requires units from multiple hospitals, fire stations, and police stations, then more decision-makers must coordinate with each other to identify the optimal way to deploy the combined rescue units.

- **Imperfect communication.** Finally, the quality of communication in such scenarios can significantly deteriorate. For instance, a disaster may damage the cell transmission towers on which rescue units rely for communication. Consequently, standard algorithms that rely on perfect communication for solving coordination problems may become impractical.

Although we use the disaster rescue scenario as the motivating domain, each of these characteristics also appears in a larger class of other multi-agent optimization (MAO) problems.

Well-known multi-agent approaches, such as distributed constraint optimization problems (DCOPs) [8, 16, 23] and the general task-allocation problem (GTAP) [24, 25], effectively address key aspects of decentralized decision-making. Both models use agents to represent decision-makers (i.e., rescue units). In DCOPs, agents assign values to their variables much as the chief of a rescue unit assigns duties to the personnel. The GTAP is similar to DCOPs in that it assigns agents to tasks. However, the GTAP introduces additional challenges imposed by realistic applications, including spatial and temporal constraints. Consequently, in this context, the rescue unit chief must create schedules for all personnel and establish ad hoc coalitions to ensure that all tasks are completed efficiently.

Both MAO problems mentioned above aim to optimize a global objective through decentralized coordination, thereby capturing the first two characteristics (i.e., global optimization and decentralized coordination). To achieve this coordination, agents communicate and coordinate their actions by exchanging messages. Unfortunately, the communication assumptions of these models are overly simplistic and often unrealistic: (1) messages are never lost, (2) all messages have very small and bounded delays, and (3) the messages sent arrive in the order that they were sent. These assumptions do not reflect the real world, where messages may be disproportionally delayed or lost due to congestion or the bandwidths of the various communication channels.

Because such MAO problems are NP-hard [16, 17], considerable research has been devoted to developing algorithms to quickly find good solutions. This class of algorithms, known as incomplete algorithms, has been the focus of numerous studies [1, 5, 6, 11, 12, 15, 17, 18, 30, 33, 34, 36, 39, 40]. Despite offering little or no quality guarantees, these algorithms have been empirically found to produce high-quality solutions. Examples of such algorithms include the distributed stochastic algorithm (DSA), maximum gain messages (MGMs), and max-sum [2, 3, 15, 39, 40] for solving DCOPs and Fisher market clearing task allocation (FMC_TA) [17] for solving the GTAP.

The presence of imperfect communication significantly affects the quality guarantees and desirable properties of such algorithms. Some of these algorithms also guarantee other desirable properties; for example, MGM-k guarantees monotonicity and convergence to a k-opt solution (i.e., a solution that cannot be improved by the action of a single agent [15, 20]). Another guaranteed property is the "anytime" property, which is achieved using the anytime framework proposed by Zivan, Okamoto, and Peled [41]. Unfortunately, algorithms and the anytime framework exploit the common simplistic communication assumptions discussed above. Consequently, the guarantees for achieving these properties may no longer hold when communication is unreliable.

The general design of most state-of-the-art MAO algorithms is synchronous. However, agents must operate in asynchronous settings [14] because the environment in which they operate is distributed, and they have no mutual clock. Therefore, synchronization is achieved by exchanging messages in each iteration of the algorithm. Specifically, in each iteration, an agent receives messages from its neighbors in the previous iteration. The agent then computes and sends messages to each of its neighbors [**ZivisanOP14**, 40]. Unfortunately, such a synchronous algorithmic design has several drawbacks when imperfect communication is considered. Due to message latency, each synchronous iteration terminates only after all messages sent in the previous iteration arrive. Therefore, advancing the algorithm from one iteration to the next depends on the longest message delay in that iteration. When messages are lost, an agent may expect to receive a message from its neighbor, whereas the neighbor is unaware that the message it sent was lost. Thus, these agents are deadlocked, with each waiting for a message from the other.

This thesis addresses communication-aware MAO (CA-MAO) problems by developing incomplete algorithms that adapt to varying communication reliabilities. Specifically, we focus on two approaches that represent such problems: DCOPs, which provide an abstract representation of multi-agent coordination, and the GTAP, which addresses a broader range of challenges typically encountered in real-life scenarios. The main tasks of this thesis are (1) model design, which involves creating communication-aware MAO models that represent dynamic communication uncertainties (i.e., CA-DCOP and CA-GTAP); (2) algorithm design, in which communication-aware algorithms are designed to address communication characteristics; and (3) analytical and empirical evaluations, which assess the efficiency of the algorithms vis-á-vis computational requirements, communication costs, solution quality, and convergence.

Chapter 2 provides the foundation for the DCOP and GTAP models, detailing their background and relevant information. In addition, it outlines the benchmark algorithms and methods used for problem solving and related work on imperfect communication in multi-agent systems.

Chapter 3 addresses the limitations of DCOPs given imperfect communication by presenting the communication-aware DCOP model (CA-DCOP), investigating how message latency and loss affect synchronous benchmark local search algorithms, and proposing an asynchronous approach for these algorithms. In addition, we develop an asynchronous anytime framework that allows for the best solution explored in nonmonotonic asynchronous local search DCOP algorithms.

Chapter 4 examines the susceptibility of message latency to the properties ensured by the MGM-k algorithm. We introduce latency-aware monotonic distributed local search (LAMDLS) and LAMDLS-2 algorithms, which guarantee monotonicity and 1- and 2-opt solutions, respectively, upon convergence.

Chapter 5 examines inference algorithms in the context of CA-DCOPs. This study comprehensively analyzes both the synchronous and alternative asynchronous designs, considering imperfect communication in the max-sum algorithm and its variants.

Chapter 6 evaluates how message latency and loss affect a real-world domain (i.e., the GTAP) by introducing the communication-aware GTAP model (CA-GTAP). We focus on the FMC_TA algorithm and address its limitations. We propose FMC_ATA, an asynchronous version of FMC_TA that is robust against message latency and loss and is more suitable in such scenarios. Finally, we investigate the conditions under which the distributed version of the algorithm should be preferred to the centralized version.

Finally, Chapter 7 summarizes this dissertation, lists its contributions, and proposes future work.

Chapter 2

Background

This chapter provides details of existing MAO models and algorithms that presume perfect communication, specifically looking at the two models DCOP and GTAP. In addition, it reviews previous research on communication awareness in MASs and the methodologies used to evaluate the performance of algorithms in scenarios with imperfect communication.

We delve into the classification of DCOP algorithms, with a particular focus on local search and inference as solution strategies [9]. In the context of local search algorithms, we consider the DSA [40] and its variant, DSA with slope-dependent probability (DSA-SDP), which moderately improves DSA's exploratory capabilities [41]. Furthermore, we address an existing anytime mechanism that can be used in conjunction with distributed synchronous local search algorithms to maintain a record of the best solutions [41].

Subsequently, we investigate the k-opt and region-opt algorithms [19, 35], focusing in particular on (MGM-k) [15] and distributed asynchronous local optimization-k (DALO-k). These algorithms guarantee that the solutions obtained are k-opt [13], implying that they cannot be improved by a group of k agents. MGM-k is a synchronous algorithm, whereas DALO-k is asynchronous.

For inference algorithms, we focus on adaptations of belief propagation [21, 38] for DCOPs [2, 3]. These include the max-sum algorithm and its variants, such as damped max-sum and max-sum with split constraint factor graph [3]. In addition, we provide details of the backtrack cost tree to evaluate the performance of these algorithms [42].

For the GTAP model (an abstraction of the law enforcement problem, which serves as a specific instance of GTAP, as detailed by Nelke, Okamoto, and Zivan [17]), we shift our focus to methodologies to deal with task allocation and optimization within complex real-world scenarios. We provide details of both the static and dynamic versions of the model and introduce the Fisher market-clearing task-allocation (FMC_TA) algorithm, as proposed by Nelke, Okamoto, and Zivan [17], highlighting its effectiveness in addressing such problems.

Chapter 3

Communication-Aware Local Search for Distributed Constraint Optimization

The synchronous design significantly affects the performance of distributed local-search algorithms for DCOPs in the face of imperfect communication. Given message latency, each synchronous iteration is completed only after the arrival of all messages sent in the previous iteration, so the advancement of the algorithm from one iteration to the next depends on the longest message delay in that iteration. When messages can be lost, an agent may expect to receive a message from its neighbor, while the neighbor is not aware that the message it sent did not arrive. Thus, these agents are deadlocked, each waiting for the message from the other.

This chapter makes the following contributions to the literature:[1]

1. We propose *communication-aware DCOPs* (CA-DCOPs),[2] which extend the DCOP model and in which patterns of communication disturbances (e.g., message latency and message loss) can be represented.

2. We demonstrate that existing distributed local search DCOP algorithms are not robust against imperfect communication. Thus, we analyze the performance and properties of standard local search algorithms after they are adjusted to perform asynchronously in scenarios that include message latency and loss.

3. We propose an asynchronous anytime mechanism that allows any local search algorithm running in an environment with imperfect communication to report the best solution it generates in its run.

[1]This chapter is based on our papers published in the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021) [29] and its extended version published in the Journal of Artificial Intelligence Research (JAIR) [27].

[2]When referring to communication awareness, we do not mean that agents are aware of the communication pattern but rather that the algorithms are designed to overcome communication limitations.

4. We show that the presence of imperfect communication can positively affect exploitative asynchronous local search algorithms. Our empirical results reveal that solution quality may improve as the quality of communication degrades (in terms of both message latency and message loss).

Our analysis and evaluation indicate that imperfect communication generates a discrepancy between the knowledge that agents hold and the actual state of the system. Thus, an agent may perform an action that is exploitative with respect to the information it holds, expecting to improve its state and the global state. However, in reality, its action degrades its state, and possibly the global state, and it unknowingly explores an unexpected part of the search space. Such explorative actions often expose agents to higher-quality solutions, allowing them to converge to better solutions.

Chapter 4

Latency-Aware k-Opt Monotonic Local Search for Distributed Constraint Optimization

Considering how message latency can negatively affect the performance of distributed algorithms, we examined a class of local-search DCOP algorithms. These algorithms are designed to provide k-opt guarantees, thereby ensuring that the solutions they generate cannot be improved by a group of k agents [19]. One such example is the MGM algorithm, a 1-opt algorithm, and its extension, MGM-2, a 2-opt algorithm. Unfortunately, their synchronous designs exploit the overly simplistic communication assumptions in the DCOP model, which do not reflect real-world scenarios. To address these limitations, researchers introduced an asynchronous k-opt algorithm called DALO to solve DCOPs. However, as we show in this chapter, DALO's design lacks robustness in scenarios with message delays, restricting its applicability.

This chapter proposes latency-aware algorithms that guarantee 1-opt and 2-opt solutions and are robust against message latency. [1]

1. We introduce latency-aware monotonic distributed local search (LAMDLS), which is a local search algorithm that is resilient against message latency. It is guaranteed to be monotonic and converges to a 1-opt solution, which is similar to the properties of the MGM algorithm mentioned by Maheswaran, Pearce, and Tambe [15]. The algorithm uses an ordered coloring scheme to prevent neighboring agents from simultaneously replacing assignments and preventing agents from waiting for messages, which is a feature of MGM algorithms.

[1] This chapter is based on our novel algorithms, which are published in the Journal of Artificial Intelligence Research (JAIR) [27] and are currently being developed further for submission to the 33rd International Joint Conference on Artificial Intelligence (IJCAI 2024) [28].

2. We present an extended version of LAMDLS, called LAMDLS-2, which allows agents to collaborate in pairs and coordinate the selection of their value assignments while preserving monotonicity. The result is a 2-opt solution. LAMDLS-2 facilitates sequential adjustments of the values among partnering agents. Using a unique pairing-selection process and an ordered coloring scheme, agents can concurrently modify the values for unconstrained pairs.

3. This work includes proofs of the theoretical properties and empirical evaluation of the proposed algorithms. Specifically, we present theoretical evidence supporting the monotonicity and convergence of the 1- and 2-opt solutions for LAMDLS and LAMDLS-2, respectively. To validate these findings, we empirically evaluate them in diverse environments with varying latency patterns. We compare the performance of our algorithms with that of MGM and MGM-2, and the results demonstrate that LAMDLS and LAMDLS-2 converge significantly faster.

Chapter 5

The Effect of Asynchronous Execution and Imperfect Communication on Max-Sum Belief Propagation

Chapter 3 investigates how message latency and loss affect standard distributed local search algorithms (e.g., MGM and DSA) and shows that imperfect communication has a significant positive effect on the performance of the asynchronous versions of these algorithms [27, 29]. Imperfect communication generates an exploration effect that significantly improves the quality of the solutions found. The present chapter focuses on investigating how imperfect communication affects distributed incomplete inference algorithms (e.g., max-sum), which have been very successful [2, 3].

Max-sum has been presented as both an asynchronous and synchronous algorithm [4, 6, 43]. To the best of our knowledge, no studies have yet focused on how these different executions of the algorithm affect its performance. Moreover, when message loss is considered, the synchronous version is not applicable because an agent may remain idle while waiting for a lost message. Although in the synchronous version message latency does not affect the actions that agents perform (it only delays them), it is intuitively expected to strongly affect the performance of the asynchronous version because the beliefs included in the messages are used by agents to construct the beliefs that they propagate to others and in their assignment selection. In asynchronous execution, belief construction and assignment selection may be performed while considering imbalanced and inconsistent information.

This chapter makes the following contributions:[1]

[1]This chapter is based on our papers published in the 27th International Conference on Principles and Practice of Constraint Programming (CP 2021) [45] and its extended version published in the Artificial Intelligence Journal (AIJ) [44].

1. We analyze the properties of the two execution versions of the max-sum algorithm, synchronous and asynchronous. More specifically, using backtrack cost trees [42], we investigate the possible differences between propagated beliefs in synchronous and asynchronous executions of max-sum.

2. We investigate how damping affects the asynchronous max-sum algorithm. Although clear evidence exists (both empirical and theoretical) that damping improves the performance of the synchronous version of the max-sum algorithm [3, 42], to the best of our knowledge, no research has yet focused on how damping affects the asynchronous version of the max-sum algorithm. We therefore analyze this effect both theoretically and empirically. Both approaches indicate that damping reduces the differences between synchronous and asynchronous executions of the max-sum algorithm.

3. We investigate how different versions of the max-sum algorithm perform in the presence of message latency and message loss. Although the beliefs propagated and the computations performed by agents are not affected by message latency in the synchronous version of the algorithm (only delayed), such is not the case for the asynchronous version of the algorithm. The empirical results reveal that damping reduces these differences. Moreover, the version of max-sum proposed by Cohen, Galiki, and Zivan [3], which includes both damping and splitting, maintains its fast convergence and solution quality, even in asynchronous execution with message delays and numerous lost messages.

Chapter 6

Asynchronous Communication Aware Multi-Agent Task Allocation

This chapter focuses on task-allocation problems, which present significant challenges in realistic scenarios. These challenges extend beyond the coordination of agents with degradation and unreliable communication. Furthermore, such scenarios are often highly dynamic and characterized by the emergence of new events or changes in the status of existing events [37]. The identification of dynamic events is most likely due to the agents in the environment. Thus, we expect agents to reinitialize the solving process when necessary [7, 9].

The FMC_TA algorithm [17] was proposed to solve problems involving a team of heterogeneous agents who need to cooperate in an environment that includes multiple tasks requiring agents with different skills ,for more details see ... in the background chapter. FMC_TA dominates state-of-the-art centralized and distributed task-allocation algorithms, including general optimization algorithms such as simulated annealing and designated algorithms such as coalition formation with a look ahead [6, 17, 31].

However, similar to the observations made in previous chapters for DCOP algorithms, FMC_TA as proposed by Nelke, Okamoto, and Zivan [17] is synchronous [15, 40, 41]. Moreover, the algorithm presented by Nelke, Okamoto, and Zivan [17] depends on agent assumptions regarding the existence and importance of the tasks to be performed. Thus, the team of agents is not independent and must rely on updates from a centralized system.

To overcome these limitations of FMC_TA when facing realistic dynamic scenarios and scenarios with imperfect communication, we follow the approach presented in previous chapters and make the following contributions:[1]

[1]This chapter is based on our paper published in the 32nd International Joint Conference on Artificial Intelligence (IJCAI 2023) [24]. Its extended version is currently being developed for submission to the Artificial Intelligence Journal (AIJ) [26].

1. We propose the Fisher market-clearing asynchronous task-allocation (FMC_ATA) algorithm, an asynchronous version of the algorithm designed to consider the possibility that messages can be delayed or lost (i.e., communication-aware design). In FMC_ATA, agents undertake computations upon receiving a message, considering the latest message from each of their neighbors. The algorithm is executed in a single distributed asynchronous phase in which agents determine both task allocation and task scheduling. Moreover, in FMC_ATA, agents can detect dynamic events (e.g., new tasks to be performed or a change in the importance of a task currently being performed) that may trigger the algorithm's execution. This capacity allows the team to adapt to the evolved scenario.

2. We demonstrate empirically that the asynchronous version (FMC_ATA) not only converges to the same solution as FMC_TA but is also robust, to some extent, against message delays and message loss.

3. We further investigate the properties of scenarios in which a distributed implementation of FMC_ATA is preferred over a centralized implementation. In the latter case, a central system is updated by the agents regarding dynamic events, calculates and updates allocation, and then updates the agents. The results show that, given message latency, a clear threshold exists before distributed performance is motivated. When message loss cannot be avoided, distributed performance is preferred.

Chapter 7

Summary and Conclusions

7.1 Conclusions

The realm of MAO faces a significant challenge in addressing the difficulties posed by imperfect communication environments. These difficulties must be overcome to represent real-world characteristics in practical applications and include message latency and message loss, which can severely affect the performance, reliability, and applicability of distributed algorithms. Our analysis identifies major limitations in the empirical and theoretical properties of such algorithms. Current MAO methods often assume ideal communication conditions, which are rarely present in real-world applications. This challenge emphasizes the need for a robust and adaptable algorithmic design that can function within the constraints of imperfect communication environments. The objective of this thesis is to address this notable gap in the existing literature by introducing communication-aware models and novel algorithmic designs. This approach produces a comprehensive suite of solutions that not only advances the field of MAO but also significantly increases the applicability of MAO algorithms in the real world, where imperfect communication is an inherent challenge.

This thesis focuses on DCOPs and GTAPs, offering a dual perspective covering both abstract and practical problems. This approach is beneficial because it enables us to thoroughly examine our hypotheses in a controlled, abstract setting and allows us to confirm our results in more practical, real-world contexts.

Chapter 3 explores how the challenges and implications of message latency and loss affect distributed local search algorithms for DCOPs and highlights the limitations of synchronous algorithms. Our investigation reveals that asynchronous versions of the DSA and MGM algorithms offer improved exploration and faster convergence. Given these promising results, we design an asynchronous anytime mechanism to mitigate the limitations of existing anytime mechanisms, which are only applicable to synchronous algorithms, as proposed by Zivan, Okamoto, and Peled [41]. Although asynchronous versions improve the results in the presence of imperfect communication, they lose certain theoretical properties of their synchronous counterparts, necessitating the development of novel solutions, which are the focus of Chapter 4.

Chapter 4 presents the LAMDLS algorithm for resolving DCOPs and ensuring convergence to 1-opt solutions. Next, we extend our design to LAMDLS-2, which guarantees convergence to 2-opt solutions. The results reveal that our approach, based on the ordered color scheme, enables agents to compute their assignments more proactively, thereby reducing the coordination effort required and significantly accelerating the convergence compared with MGM and MGM-2.

In this continuation of our investigation, we delve into inference algorithms, focusing specifically on the max-sum algorithm and its variants. Our analysis examines the differences between synchronous and asynchronous executions and the benefits of using damping to mitigate the adverse effects of message latency. Our theoretical and empirical examinations enhance our understanding of the max-sum algorithm in distributed settings where communication challenges are prevalent.

The research also focuses on the design of the FMC_ATA algorithm, which is a robust solution for multi-agent task-allocation problems that provides high-quality results, even under adverse communication conditions. This algorithm is a significant improvement over its predecessors in terms of real-world applicability. Furthermore, we compare the performance of distributed algorithms with that of centralizing all information and applying central solving. Finally, the results establish the conditions under which distributed algorithms are justified.

In conclusion, this thesis challenges the essential difficulties created by imperfect communication in MAO by generating creative and adaptable solutions. By developing new algorithms, this work fills a crucial gap in existing research and substantially improves the practical application of MAO. These results and analyses establish a strong foundation for future advancements in MAO, which should inspire further research to create more resilient and efficient systems.

7.2 Future Work

This thesis establishes a solid foundation for exploring numerous potential directions for the development of practical MAO algorithms. A promising area for future exploration is the additional constrained communication graphs that reflect real-world challenges, such as limited bandwidth or unique network topologies, exemplified by the smart home problem identified by Rust, Picard, and Ramparany [32], where algorithms reduce the communication costs of agents. The asynchronous algorithmic strategies developed in this study, designed with communication awareness in mind, may face limitations in these contexts.

This study reveals a gap in the comprehension of the practical benefits provided by asynchronous variants of distributed algorithms in the context of imperfect communication. These algorithms either maintain the accuracy of the results with respect to their synchronous counterparts or even improve them. A potential avenue for future research involves applying explainable artificial intelligence techniques to improve our understanding of the underlying mechanisms that allow agents to operate effectively with outdated information in asynchronous settings.

Finally, future research may also focus on refining the LAMDLS and LAMDLS-2 algorithms introduced in this study. Building on our initial contributions, we can extend our methodology to include a general k-opt algorithm. Exploring the DALO-k algorithm, as discussed by Kiekintveld et al. [13], is another valuable avenue of research. Adapting its design to include communication-aware features could help address some of the limitations.

Bibliography

[1] M. Basharu, I. Arana, and H. Ahriz. “Solving DisCSPs with penalty driven search”. In: *Proceedings of AAAI*. 2005, pp. 47–52.

[2] Ziyu Chen et al. “A class of iterative refined Max-sum algorithms via non-consecutive value propagation strategies”. In: *Auton. Agents Multi Agent Syst.* 32.6 (2018), pp. 822–860.

[3] Liel Cohen, Rotem Galiki, and Roie Zivan. “Governing convergence of Max-sum on DCOPs through damping and splitting”. In: *Artificial Intelligence Journal (AIJ)* 279 (2020).

[4] Yanchen Deng and Bo An. “Speeding Up Incomplete GDL-based Algorithms for Multi-agent Optimization with Dense Local Utilities”. In: *Proceedings of the 29th International Joint Conference on Artificial Intelligence, (IJCAI)*. 2020, pp. 31–38.

[5] A. Farinelli et al. “Decentralised Coordination of Low-Power Embedded Devices Using the Max-Sum Algorithm”. In: *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*. 2008, pp. 639–646.

[6] A. Farinelli et al. “Decentralized Coordination of Low-Power Embedded Devices Using the Max-Sum Algorithm”. In: *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*. 2008, pp. 639–646.

[7] Alessandro Farinelli, Luca Iocchi, and Daniele Nardi. “Distributed on-line dynamic task assignment for multi-robot patrolling”. In: *Autonomous Robots* 41.6 (2017), pp. 1321–1345.

[8] F. Fioretto, E. Pontelli, and W. Yeoh. “Distributed Constraint Optimization Problems and Applications: A Survey”. In: *Journal of Artificial Intelligence Research* 61 (2018), pp. 623–698.

[9] Ferdinando Fioretto, Enrico Pontelli, and William Yeoh. “Distributed constraint optimization problems and applications: A survey”. In: *Journal of Artificial Intelligence Research* 61 (2018), pp. 623–698.

[10] Ferdinando Fioretto, William Yeoh, and Enrico Pontelli. "A Multiagent System Approach to Scheduling Devices in Smart Homes". In: *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*. 2017, pp. 981–989.

[11] K. D. Hoang et al. "A Large Neighboring Search Schema for Multi-agent Optimization". In: *Proceedings of CP*. 2018, pp. 688–706.

[12] E. G. Jones, M. B. Dias, and A. Stentz. "Learning-enhanced Market-based Task Allocation for Oversubscribed Domains". In: *Intelligent Robots and Systems, 2007. IROS 2007. IEEE/RSJ International Conference on*. IEEE. San Diego, CA, 2007, pp. 2308–2313.

[13] Christopher Kiekintveld et al. "Asynchronous algorithms for approximate distributed constraint optimization with quality bounds." In: *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. Vol. 10. 2010, pp. 133–140.

[14] N. A. Lynch. *Distributed Algorithms*. Morgan Kaufmann Series, 1997.

[15] R. Maheswaran, J. Pearce, and M. Tambe. "Distributed Algorithms for DCOP: A Graphical Game-Based Approach". In: *Proceedings of PDCS*. 2004, pp. 432–439.

[16] P. J. Modi et al. "ADOPT: Asynchronous Distributed Constraint Optimization with Quality Guarantees". In: *Artificial Intelligence* 161.1–2 (2005), pp. 149–180.

[17] Sofia Amador Nelke, Steven Okamoto, and Roie Zivan. "Market Clearing-based Dynamic Multi-agent Task Allocation". In: *ACM Transactions of Intelligent Systems Technology.* 11.1 (2020), 4:1–4:25.

[18] D. T. Nguyen et al. "Distributed Gibbs: A Linear-Space Sampling-Based DCOP Algorithm". In: *Journal of Artificial Intelligence Research* 64 (2019), pp. 705–748.

[19] J. Pearce and M. Tambe. "Quality Guarantees on k-Optimal Solutions for Distributed Constraint Optimization Problems". In: *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*. 2007, pp. 1446–1451.

[20] Jonathan P Pearce and Milind Tambe. "Quality Guarantees on k-Optimal Solutions for Distributed Constraint Optimization Problems". In: *Proceedings of IJCAI*. 2007, pp. 1446–1451.

[21] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Francisco, California: Morgan Kaufmann, 1988.

[22] Arseni Pertzovskiy, Roie Zivan, and Noa Agmon. "CAMS: Collision Avoiding Max-Sum for Mobile Sensor Teams". In: *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*. 2023, pp. 104–112.

[23] A. Petcu and B. Faltings. "A Scalable Method for Multiagent Constraint Optimization". In: *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI)*. 2005, pp. 1413–1420.

[24] Ben Rachmut, Sofia Amador Nelke, and Roie Zivan. "Asynchronous Communication Aware Multi-Agent Task Allocation". In: *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence (IJCAI-23)*. 2023, pp. 262–270.

[25] Ben Rachmut, Sofia Amador Nelke, and Roie Zivan. "Asynchronous Communication Aware Multi-Agent Task Allocation (EA)". In: *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*. 2023, pp. 2340–2342.

[26] Ben Rachmut, Sofia Amador Nelke, and Roie Zivan. "Asynchronous Communication Aware Multi-Agent Task Allocation[in progress]". In: *Artificial Intelligence* (2024).

[27] Ben Rachmut, Roie Zivan, and William Yeoh. "Communication-Aware Local Search for Distributed Constraint Optimization". In: *Journal of Artificial Intelligence Research* 75 (2022), pp. 637–675.

[28] Ben Rachmut, Roie Zivan, and William Yeoh. "Latency-Aware 2-Opt Monotonic Local Search for Distributed Constraint Optimization [in progress]". In: *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI-24)*. 2024.

[29] Ben Rachmut, Roie Zivan, and William Yeoh. "Latency-aware local search for distributed constraint optimization". In: *Proceedings of the 2021 International Conference on Autonomous Agents and Multiagent Systems*. 2021, pp. 1019–1027.

[30] S. D. Ramchurn et al. "Coalition formation with spatial and temporal constraints". In: *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-10)*. Toronto, Canada, 2010, pp. 1181–1188.

[31] Sarvapali D. Ramchurn et al. "Decentralized Coordination in RoboCup Rescue". In: *Computer* 53.9 (2010), pp. 1447–1461.

[32] Pierre Rust, Gauthier Picard, and Fano Ramparany. "Resilient distributed constraint reasoning to autonomously configure and adapt IoT environments". In: *ACM Transactions on Internet Technology* 22.4 (2022), pp. 1–31.

[33] A. Schoneveld, J. F. de Ronde, and P. M. A. Sloot. "On the Complexity of Task Allocation". In: *Journal of Complexity* 3 (1997), pp. 52–60.

[34] M. Smith and R. Mailler. “Getting What You Pay For: Is Exploration in Distributed Hill Climbing Really Worth It?” In: *Proceedings of IAT*. 2010, pp. 319–326.

[35] M. Vinyals et al. “Quality Guarantees for Region Optimal DCOP algorithms”. In: *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 2011, pp. 133–140.

[36] W.E. Walsh and M.P. Wellman. “A market protocol for decentralized task allocation”. In: *Proceedings of the International Conference on Multi-Agent Systems*. 1998, pp. 325 –332.

[37] Changyun Wei, Koen V Hindriks, and Catholijn M Jonker. “Dynamic task allocation for multi-robot search and retrieval tasks”. In: *Applied Intelligence* 45.2 (2016), pp. 383–401.

[38] Chen Yanover, Talya Meltzer, and Yair Weiss. “Linear Programming Relaxations and Belief Propagation - An Empirical Study”. In: *Journal of Machine Learning Research* 7 (2006), pp. 1887–1907.

[39] M. Yokoo and K. Hirayama. “Distributed breakout algorithm for solving distributed constraint satisfaction problems”. In: *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*. 1996.

[40] W. Zhang et al. “Distributed Stochastic Search and Distributed Breakout: Properties, Comparison and Applications to Constraint Optimization Problems in Sensor Networks”. In: *Artificial Intelligence* 161.1–2 (2005), pp. 55–87.

[41] R. Zivan, S. Okamoto, and H. Peled. “Explorative anytime local search for distributed constraint optimization”. In: *Artificial Intelligence* 211 (2014).

[42] Roie Zivan, Omer Lev, and Rotem Galiki. “Beyond Trees: Analysis and Convergence of Belief Propagation in Graphs with Multiple Cycles”. In: *Proceedings of the 34th International Conference of the Association for the Advancement of Artificial Intelligence (AAAI)*. 2020, pp. 7333–7340.

[43] Roie Zivan et al. “Balancing exploration and exploitation in incomplete Min/Max-sum inference for distributed constraint optimization”. In: *Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS)* 31.5 (2017), pp. 1165–1207.

[44] Roie Zivan et al. “Effect of asynchronous execution and imperfect communication on max-sum belief propagation”. In: *Autonomous Agents and Multi-Agent Systems* 37.2 (2023). ISSN: 1387-2532.

[45] Roie Zivan et al. “The effect of asynchronous execution and message latency on Max-Sum”. In: *27th International Conference on Principles and Practice of Constraint Programming (CP 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik. 2021.