Consensus-Based Auction Methods with Bid Intercession for SAR

Master of Science Thesis

Victor Guillet

Delft University of Technology

RÉPUBLIQUE FRANÇAISE
Liberté
Égalité
Fraternité

ONERA
THE FRENCH AEROSPACE LAB

Consensus-Based Auction Methods with Bid Intercession for SAR

by

Victor Guillet

to obtain the degree of Master of Science

at the Delft University of Technology,

to be defended publicly on Tuesday December 19th, 2023.

Student number:	4488636
Project duration:	January 2023 – December 2023
Thesis committee:	Dr. O.A. Sharpanskykh
	Dr. D.M. Pool
	Dr. B.F. Santos
	Dr. C. Lesire

Acknowledgements

This master thesis concludes my studies at the Delft University of Technology. These years have been the most interesting and eye-opening of my life so far, and I consider myself fortunate to have bonded, grown, and collaborated alongside some of the brightest and most interesting people I have encountered so far. Navigating the challenges inherent in these studies, coupled with the high standards upheld, has not only taught me the necessary skills and knowledge for an exciting career but has also profoundly shifted my perspective on life. The lectures, the projects (studies and personal), my time in DARE, and the endless great discussions with my close friends have instilled a deep fascination for the intricacies of the world around us. Furthermore, I consider myself incredibly lucky to have in time joined the ONERA, the French Aerospace Labs. It in turn has taken on the role of continuing the task started by TU Delft in further expanding the horizon, and providing a depth to the world which never ceases to amaze and thrill me.

Consequently, I would like first to express my gratitude to all of my supervisors (in no particular order): Dr. Alexei Sharpanskykh, Dr. Charles Lesire, Dr. Gauthier Picard, and Dr. Christophe Grand. Dr. Alexei Sharpanskykh provided invaluable insights into methodology, processes, and significantly enriched my academic education. Dr. Charles Lesire offered meticulous guidance in research methodologies, aiding me in navigating the research landscape to properly frame the core topics and find focus. Dr. Gauthier Picard's profound expertise in the field and guidance on the research's theoretical aspects were instrumental and key in reaching the conclusions and results presented here. Finally, Dr. Christophe Grand's emphasis on practical applications significantly enriched the research's pragmatic dimensions, ensuring results grounded in meaningful real-world contexts. Their collective expertise and continuous support have been pivotal in shaping the trajectory of this research and my growth as a master student.

Secondly, thank you to all my friends, colleagues and teachers, who shared the experiences with and who rode with me alongside every step of the journey. I want to thank in particular (in no specific order) Arthur Thiam, Mihaly Katona, Xavier Goby, Julien Svenson, the french and Dubai crews, and all my other friends made along the way for the great times and discussions. You truly made this time in Delft special, and I consider the shared moments to have been pivotal in shaping those memorable years.

And finally, I would like to thank all my family. I can say without hesitation that they truly are the most inspirational figures in my life. Their unconditional love, support, and guidance not only enabled me to complete my long studies but have made me who I am today. Thank you for the patience, care, and trust you have given me, I would not be here today without you.

This master thesis concludes my time in TU Delft and marks the beginning of the next step of this fantastic journey I have embarked on. So with this, let's go see what the future holds!

Victor Guillet
Delft, December 2023

Contents

List of Figures

List of Tables

List of Abbreviations

List of Abbreviations

ABM Agent-Based Modeling

ACBBA Asynchronous Consensus-Based Bundle Algorithm

BRS Beta Reputation System

CBAA Consensus-Based Auction Algorithm

CBBA Consensus-Based Bundle Algorithm

CCBBA Coupled-Constraint Consensus-Based Bundle Algorithm

CD Complex Dependencies

D Dynamic tasks and environment

DRS Dirichlet Reputation System

G Global Communication

IA Instantaneous knowledge available

ID In-schedule Dependencies

K Known a priori knowledge

L Local Communication

MBA Market Based Approach

MILP Mixed Integer Linear Programming

MR Multiple robot

MRTA Multi-Robot Task Allocation

MT Multiple task

N No comms uncertainty

NBR Naive Binary Rating

ND No Dependencies

PA Positive Asymmetry

PDF Probability Density Function

POI(s) Point Of Interest

PSI Parallel Single Item Auction

PSO Particle Swarm Optimisation

RCRS RoboCup Rescue Simulation

S Static tasks and environment (or) Stochastic comms

SAR Search And Rescue

SI Single Item Auction

SR Single robot

SSI Sequential Single Item Auction

ST Single Task

TA Time-extended knowledge available

TA:SP Synchronisation constraints with other tasks

TA:TW Time-window constraints

TCBBA Team Consensus-Based Bundle Algorithm

U Unknown a priori knowledge

WDP Winner Determination Process

XD Cross-schedule Dependent

Introduction

This document reports on the ten months of research undertaken as part of my Master of Science degree in Aerospace Engineering at the Delft University of Technology.

This research aims to propose a novel approach to tackling **task allocation problems** akin to those encountered in **Search And Rescue** (SAR) challenges, the main case study of this paper. Task allocation refers to the process of effectively assigning various tasks or activities to specific agents, or entities within a system to achieve predefined objectives or goals efficiently. In the case of SAR example, firefighters, ambulances, and police forces must coordinate to deal with the aftermath of an earthquake by controlling and extinguishing fires, rescuing victims, and clearing roads. Such systems therefore prove critical to ensuring an effective and reliable response in high-stake applications, where speed and efficiency are key to enabling a successful outcome. Additionally, the unpredictable nature of disasters and their consequent impact leads to the imperative need for a coordination protocol capable of operating reliably under any circumstance, given an effective response often means the difference between life and death for many.

This research initiative originated from an observation made during my master's internship: The absence of a universal solution for task allocation forces designers to choose algorithms based on specific application constraints. Particularly, the reliability of communication among agents within a fleet, such as drones or robots, emerged as a primary consideration. At the time of this writing, two major categories of allocation methods can be observed; **Methods assuming reliable communications** (which are therefore unable to cope with communication failures but often result in more optimal solutions), and **Methods designed for unreliable communications** (which are more robust to communication failure at the expanse of optimality). This consequently results in a significant tradeoff of optimality for the sake of robustness.

In an attempt to address this, the research proposes a novel highly customisable approach capable of seamlessly combining multiple existing methods together to capitalise on existing resources effectively. It additionally significantly expands the programmability and control over such processes, allowing for structuring and organising a coordination process beyond what was possible until now while still retaining critical performance and robustness guarantees critical to high-stake applications. The resulting solution therefore unlocks an entire spectrum of hybrid approaches and opens exciting avenues for novel task allocation strategies. The research therefore primarily focuses on analysing the core properties of this solution, its subsequent applications, and emerging dynamics with the intent to pave the way for future studies.

This project was conducted in collaboration with the ONERA, the French Aerospace Center, where I spent six months in Toulouse, France. This immersion provided direct access to cutting-edge resources and expertise, which I believe amplified the depth of this work. Collaborating with experts and peers sparked creative discussions, while exposure to real-world projects deepened my understanding of the subject and its applications. This experience significantly shaped and enhanced the project's scope and outcomes.

This thesis report is organised as follows: In Part I, the scientific paper is presented, and Part II contains the relevant Literature Study that supports the research.

I

Scientific Paper

Consensus-Based Auction Methods with Bid Intercession for SAR

Victor Guillet*

Delft University of Technology, Delft, The Netherlands

Abstract

This research addresses the process of task allocation in a heterogeneous multi-agent fleet through the introduction of a novel mechanism in existing decentralised consensus algorithms: bid intercession. Bid intercession refers to the principle of agents biding on behalf of other agents in decision-making architectures leveraging market-based decision strategies. The method exploits and extends existing consensus-based allocation processes through the redistribution of responsibilities in the auction process to achieve various degrees of centralisation in the task allocation process. It is demonstrated that the extension proposed allows for hybridising multiple allocation methods together and structuring the auction process (notably through embedding hierarchies and decision trees directly in the decision-making process) all the while retaining the convergence robustness and performance guarantees provided by the underlying algorithms. The Search and Rescue case study is investigated to assist in framing the research and provide a reference scenario for the application of such concepts. This concept, unexplored so far in consensus-based approaches, not only opens up a sway of coordination architectures and optimisations but also paves the way to novel ethically compliant autonomous systems while retaining essential performance and robustness properties crucial in high-stake applications.

Keywords: Task Allocation, Distributed Robot Systems, Multi-Robot Systems, Auctions, Consensus

1 Introduction

The usage of heterogeneous robot fleets has grown significantly in recent years across a number of industries and sectors, such as agriculture [1, 2], logistics [3], and warehouse management [4]. As their capability develops, it is becoming clear that these systems hold a lot of potential in an expanding range of applications. However, higher autonomy and intelligence are required to truly unlock new use cases. One such area is the robots' communication and coordination mechanisms, which strive to best coordinate a collection of agents and the various resources available within a network to take on a number of tasks.

Autonomous task allocation, also known as the *Multi-Robot Task Allocation problem (MRTA)* [5], has proven to be a difficult challenge to tackle, and numerous strategies have been proposed in an attempt to address it [6]. A universal approach has not yet been established, and the majority of solutions that have been suggested so far were created expressly to deal with problems of a particular nature. This directly results from the fact that the various scenarios of applications often present significant differences, both in goals and operational conditions.

Among the existing methods investigated, *auction-based* approaches have demonstrated great promise, and have been adapted to tackle a large variety of MRTA scenarios [7, 8, 9]). These techniques employ virtual bidding and auction procedures to distribute tasks within a fleet and reach a consensus on the task distribution best suited to specific circumstances. Two main categories can be established [6]; the centralised methods, concentrating the decision-making process around a single agent, and the distributed methods, which determine an allocation through peer-to-peer protocols and exchanges of information. The centralised methods generally tend to yield superior solutions but fail to cope with communication failures, while the distributed nature of decentralised approaches ensures that these protocols are capable of coping in such situations at the expense of some optimality.

This research seeks to explore the impact of the introduction of a novel technique in the process of decentralised auction-based task allocation in an attempt to address the various limitations of the original approaches. The method proposed modifies the responsibility of the various actors in an allocation network in order to blend together multiple algorithms and enable novel dynamics for structuring a consensus process. This is achieved by introducing the capability for agents to not just bid for themselves, but also *on behalf of others* in the auction processes, a mechanism unexplored until now in the literature. This work focuses specifically on the technical implementation, viability, and consequences of such a mechanism on the underlying methods used as foundation.

*Msc Student, Sustainable Air Transport, Faculty of Aerospace Engineering, Delft University of Technology

To guide this research and help frame the problem addressed, a MRTA case study was selected. We chose to focus on the *Search and Rescue (SAR)* scenario [10], which involves locating and rescuing victims in disaster-stricken areas. This scenario poses challenges due to unreliable infrastructure, evolving objectives, and partially known operating environments. Efficiently coordinating diverse agents in such scenarios is crucial for maximising resource utilisation and improving response success, as the success of the response often translates to the difference between life and death for the victims.

The research is structured as follows. First, a description of the SAR case study is provided in section 2. This section provides a detailed example of the overarching problem tackled, along with the associated challenges. The background (section 3) then contains a definition of the MRTA problem, along with a framework for analysing it, and consequent analysis of SAR. The section additionally presents the class of methods considered for tackling the case study and identifies the research gap addressed in this paper. The method proposed is then detailed in section 4, and an analysis of its properties is performed in section 5. A subsequent discussion of the method is proposed in section 7. The consequent applications and dynamics arising from the introduction of bid intercession are discussed in section 6, along with examples of applications in our SAR case study. Finally, conclusions and suggestions for future works are then provided in section 8.

2 Search And Rescue: A Case Study

The Search and Rescue (SAR) scenario refers to the challenge of performing a rescue mission in the context of a (natural or human-driven) disaster. Examples of such disasters include earthquakes, floods, areas of conflict, etc. The goal of such missions is to coordinate multiple actors (often heterogeneous to various degrees) in an attempt to locate any situations needing attention, overcome obstacles, and rescue lives and infrastructure as efficiently as possible. This task is further complexified by the uncertain and unstable nature of the operating environment, often made up of an unreliable (or entirely failing) infrastructure, adversely impacting aspects of the mission such as communication and movement. While the challenges presented by a search and rescue scenario are wide-ranging (encompassing concepts such as victim detection, autonomous navigation, and communication management), this work focuses solely on the task allocation (MRTA) aspect of the problem.

The SAR scenario proposed in the RoboCupRescue [11] challenge will be used as a reference for this study. The RoboCupRescue challenge is an international robotics competition organised by the RoboCup Foundation, an organisation founded in 1996 with the goal of promoting robotics and AI research by proposing exciting yet challenging tasks for teams to complete [11]. The extensive work already performed in validating the scenario proposed in the challenge and correspondence with real-life cases ensures a robust baseline to base ourselves on for this work. The RoboCupRescue challenge defines 6 major constraints [12] that must be accounted for :

- Limited resources
- Incomplete information
- Large number of actors involved
- Heterogeneity (in actors and tasks/task requirements)
- Real-time decision making
- Dynamic scenario state

Based on the above, the following scenario was devised by the Foundation:

2.1 Agents

Eight different agent classes are defined in the challenge. They are shown in Figure 1.

Figure 1: The 8 different agent classes are established in the RoboCup Rescue Simulator [13]

Agents can be grouped into pairs; a mobile agent and a location agent. The mobile agent is usually responsible for taking on various on-field tasks: the fire brigades extinguish the fires, the ambulances rescue the victims, the police forces clear debris from the roads, and the civilians reach shelters. Each category of agent (and possibly individual agent) displays differences in capabilities and objectives, corresponding to their respective roles and status. The coordinating (or location) agents take on the tasks of managing supplies, communication, and coordination across the whole system, and usually focus on working in tandem with their complements to ensure certain aspects of the mission are effectively handled. Furthermore, each agent category (and/or even individual agent) can also present further differences, in aspects such as (but not only) communication capabilities, or displacement speed.

2.2 Environment

The environment itself comprises roads, buildings, along with a number of key points of interest (POIs), such as fire hydrants and gas stations. The specific locations for the static agents (ambulance centre, fire station, etc) are also often defined. An example of such an environment can be seen in Figure 2.

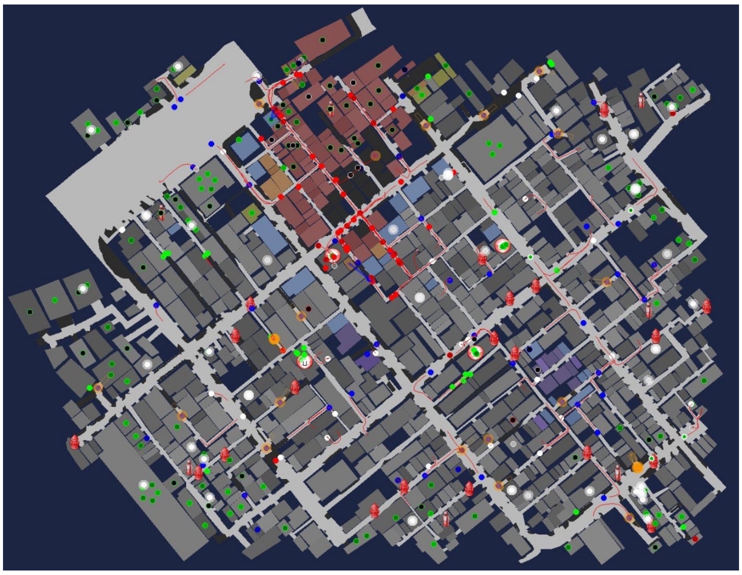

Figure 2: Example of an environment used in the RoboCupRescue competition [12]. The environment contains a large array of points of interest used to simulate a disaster scenario. The various dots represent the agents, and the colours of the building reflect their states (damaged, burning, etc).

.

2.3 Tasks

A large array of tasks can be observed in the SAR scenario. To match the above-described agents and environment setup, the following are established:

- **Victims rescue:** Those tasks are the responsibility of the ambulance teams. They must go collect victims and drop them off at an ambulance centre or shelter.
- **Fire fighting:** Those tasks are handled by the fire brigades. They must reach a specific site and control and extinguish a fire to prevent it from spreading to nearby buildings.
- **Debris clearing:** These tasks are the responsibility of the police forces, which must clear roads that have been blocked by debris to make them usable.
- **Reaching shelter:** The civilians are responsible for getting to the shelters.
- **Coordination tasks:** Those tasks may theoretically be undertaken by any agent, but will usually be the responsibility of the coordination centres. The nature of these tasks often depends on the coordination mechanisms and goals adopted in the solution.
- **Communication tasks:** Those tasks are also theoretically available to all agents. They can however be made the responsibility of more specialised agents (with better equipment, for example).
- **(Re)supply tasks:** Those tasks are available to the different agents depending on their respective needs. All mobile agents will need to resupply at their respective centres and a gas/charging station, firetrucks will need to connect to fire hydrants etc.

This MRTA problem was overall selected as it presents a particularly complex combination of constraints and objectives (more in subsection 3.2). It will be used in the rest of the research as the main context of reference for understanding the challenges addressed by our approach and highlighting consequent applications and use cases.

3 Background

In order to correctly analyse the SAR case study and devise an appropriate solution, a framework for formalising MRTA problems is first established in subsection 3.1, and applied to the SAR problem in subsection 3.2. An overview of the existing market-based methods is then provided in subsection 3.3.

3.1 Multi-Robot Assignment Problem (MRTA)

The multi-robot assignment problem, also known as the Multi-Robot Task Allocation [5] (MRTA) problem refers to the challenge of assigning N_t tasks to N_u agents, to obtain a conflict-free distribution of tasks to agents that maximises some overall reward (or minimise some overall cost). An allocation is qualified as "conflict-free" if each distinct task is assigned to at most one agent. A maximum of L_t tasks can be assigned to each agent, and the assignment is considered as completed once $N_{\min} \triangleq \min\{N_t, N_u L_t\}$ tasks have been assigned. This problem has been extensively studied and for the sake of maintaining consistency with the existing literature, the following integer (possibly non-linear) program formulation proposed in [14, section II-A] (the main paper this research builds on) will be adopted to formalise the problem.

In it, binary decision variables x_{ij} are used to represent the binary decision of whether or not a task j is assigned to an agent i:

$$\max \sum_{i=1}^{N_u} \left(\sum_{j=1}^{N_t} c_{ij}\left(\mathbf{x}_i, \mathbf{p}_i\right) x_{ij} \right) \tag{1}$$

subject to

$$\begin{aligned} \sum_{j=1}^{N_t} x_{ij} \leq L_t \quad & \forall i \in \mathcal{I} \\ \sum_{i=1}^{N_u} x_{ij} \leq 1 \quad & \forall j \in \mathcal{J} \\ \sum_{i=1}^{N_u} \sum_{j=1}^{N_i} x_{ij} = N_{\min} \triangleq & \min\{N_t, N_u L_t\} \\ x_{ij} \in \{0,1\} \quad & \forall (i,j) \in \mathcal{I} \times \mathcal{J} \end{aligned} \tag{2}$$

In the above-provided definition, decision variable $x_{ij} = 1$ if task i is assigned to agent j, and is set to 0 otherwise. $\mathbf{x}_i \in \{0,1\}^{N_t}$ is then a vector with x_{ij} as jth element. The index sets are defined as $\mathcal{I} \triangleq \{1, \ldots, N_u\}$ for the agents, and $\mathcal{J} \triangleq \{1, \ldots, N_t\}$ for the tasks. The vector $\mathbf{p}_i \in (\mathcal{J} \cup \{\emptyset\})^{L_t}$ then represents an ordered sequence of tasks for agent i; its kth element is $j \in \mathcal{J}$ if agent i conducts j at the kth point along the path, and becomes $\emptyset$ (denoting an empty task) if agent i conducts less than k tasks.

The *local reward* for agent i is therefore represented by the summation term in Equation 1. It should be noted that in this formalisation, the allocation operates around a reward function, resulting in a maximisation problem. If a cost function was instead adopted, the problem would turn into a minimisation one.

The *score function* is usually assumed to satisfy $c_{ij}\left(\mathbf{x}_i, \mathbf{p}_i\right) \geq 0$ and can be any (usually non-negative) function of either assignment $\mathbf{x}_i$ and/or path $\mathbf{p}_i$. In the case of mobile autonomous vehicles and robots, scoring/cost functions usually exploit path-dependent properties to represent the cost/reward of taking on various tasks (path length/mission completion time/etc.).

Note. *The above-described terminology shall be adopted throughout the rest of this work.*

Although a wide range of scenarios are characterised as MRTA problems, the nature of the challenges presented may vary significantly from one case to another. Understanding the precise nature of a given MRTA problem is therefore essential to comprehending the operating scenario's inherent difficulties and constraints, and in turn, finding a fitting solution. Consequently, a taxonomy was established to help understand and categorise the various aspects observed.

The taxonomy, developed over a few iterations across multiple papers, defines six main properties that must be considered when analysing a MRTA problem (more extensively described in [5, 15, 6]):

1. **Robots' capabilities:** How many tasks can a robot take up at once?
2. **Tasks types:** Can the tasks be completed by a single robot, or do they allow for/require multiple robots to cooperate/coordinate?
3. **Assignment and Constraints:** At what rate is information available, and what are the constraints inherent to the various tasks (are there time windows, a need for synchronisation, etc)?
4. **Degree of inter-dependence:** What is the degree of inter-dependence present across tasks?
5. **Communication:** This category addresses two aspects of the communication present:
   - **Connectivity Type:** Is the communication global or local (limited in space)?
   - **Uncertainty upon messages:** Is the communication reliable or lossy?
6. **Environment and Tasks:** This category qualifies the dynamics (or lack thereof) present in the operating environment. Two aspects are considered:
   - ***A priori* Knowledge:** Are the tasks known ahead of time or discovered at run time?
   - **Dynamics:** Are the tasks and environment dynamics in nature or static in time?

The full description and table of available options can be found in [6, section 3.2, tables 2,3].

Each of the categories described above allows for qualifying a specific aspect of a given MRTA instance and provides a baseline for comparing different application scenarios in a structured and organised fashion. It furthermore allows for ruling out certain methods and approaches, incompatible by nature with certain operational scenarios. One such example is the connectivity type. Lossy and local communications will often lead to ruling out centralised methods, as these approaches often prove too sensitive to communication quality. Research synthesising the existing methods and their properties may be used to hint as to which method may prove more adapted to which properties (e.g. [6, table 9 & figure 5]).

Having proposed an analysis scheme, the Search and Rescue case study may now be investigated.

3.2 Analysing MRTA within SAR

Combining the challenges and the scenario described in section 2 with the taxonomy described in subsection 3.1 allows for analysing and categorising the SAR scenario as follows:

Note. *In the event that a given scenario displays properties falling in multiple categories, it is convention to retain the most constraining one. This however does not apply when the properties do not "overlap" or encompass each other.*

1. **Robots' capabilities – MT:** While most mobile agents are limited to single tasks, the coordinating agents are able to for example handle resupply and communication tasks simultaneously.
2. **Tasks types – MR:** The tasks available during a SAR mission are varied, and range from simply picking up a victim, which would be classified as single robot (**ST**) to coordinating multiple agents to prevent the spread of a fire for example, which would be considered as multi-robot (**MR**).
3. **Assignment and Constraints – TA:TW:SP:** Tasks such as picking up a victim could depend on a road being cleared, presenting a synchronisation constraint (**SP**). A fire must be dealt with rapidly to prevent it from spreading, resulting in a time-window constraint (**TW**). Finally, a civilian calling to notify the various first-response services would be considered instantaneous knowledge (**IA**), whereas having an agent on scene continuously relaying information back to the coordination centres would be considered time-extended knowledge (**TA**).
4. **Degree of inter-dependence – CD:** The SAR problem presents tasks with both no dependencies **ND** (e.g. going to pick up a victim in an area with cleared roads), in-schedule dependencies **ID** (e.g. a firetruck must resupply first before taking on more fire-fighting tasks), cross-schedule (**XD**) (e.g. a road must first be cleared by the police force to be used to rescue a victim), and finally the choice of the ordering of tasks affect the final outcome of the mission, leading to the **CD** classification.
5. **Communication:**
   - **Connectivity Type – L:** The communication is local as the SAR scenario assumes a failing or unreliable infrastructure

- **Uncertainty upon messages - S:** The communication is assumed to be lossy/unreliable for the same reasons as above (mostly due to infrastructure).

6. **Environment and Tasks:**
   - ***A priori* Knowledge – U:** The tasks are unknown (**U**) ahead of time.
   - **Dynamics – D:** As specified by characteristic #6, the scenario and environment are dynamic in nature.

Having established the above breakdown of the problem characteristics, it is now possible to start considering various appropriate solutions for addressing this problem.

3.3 Market-Based Methods

While a number of different methods may prove interesting to consider when tackling such problems, it was decided for this work to focus on the Market-Based Approaches[1] (*MBA*s). It was chosen as such approaches are able to remain efficient while still offering a great level of flexibility and liberty in their design. Those methods come in many flavours, and importantly can be adjusted to control the level of centralisation of the coordination process (more on this and why it is important below). Additionally, they have already proven quite successful at tackling this category of problem [7, 8, 9, 16, 17].

Market-based approaches are a well-established and researched group of methods. They seek to solve the MRTA problem by providing the participating agents with a simulated economic environment in which to trade resources and tasks. Agents are provided with a reward/cost function to calculate their own utility, which is in turn leveraged to carry out the task assignment. Many variations have been developed, each seeking to strike a balance between a few crucial inherent characteristics to fit a given application scenario:

1. The process's degree of centralisation.
2. The task allocation rate.
3. The bid estimation & mechanism.

When attempting to categorise the various existing methods, two broad categories can be distinguished:

1. **Auction-based methods:** Those methods rely on the centralisation of information around an auctioneer to carry out the allocation process. These methods tend to produce the most optimal results in ideal conditions but are not very robust as the centralisation of information introduces a single point of failure. Notable examples include *Single Item Auctions* [18] (SI) (depicted in Figure 3), *Parallel Single Item Auctions* [19] (PSI), *Sequential Single Item auctions* [20] (SSI), and *Combinatorial Auctions* [21].

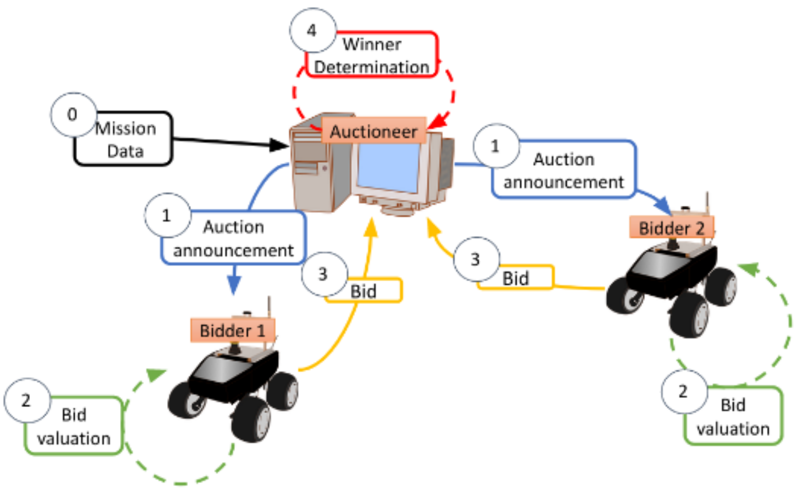

Figure 3: The Single-Item (SI) auction process [18], an example of Auction-based methods (explained in details in [6, figure 1])

2. **Consensus-Based approaches:** Those methods rely on a peer-to-peer exchange of information to perform the task allocation process. They avoid a single point of failure at the expense of some optimality. The two main examples of such approaches are *Consensus-Based Auction Algorithm* (CBAA) and *Consensus-Based Bundle*

[1] It is important to note here that we specifically focus on the underlying auction principles of market systems, and not other aspects such as pricing mechanisms, etc.

Algorithm (CBBA)[14]. An illustration of the CBBA process is provided in Figure 4 as an example. This category of approaches is described in detail in [6, section 5.2].

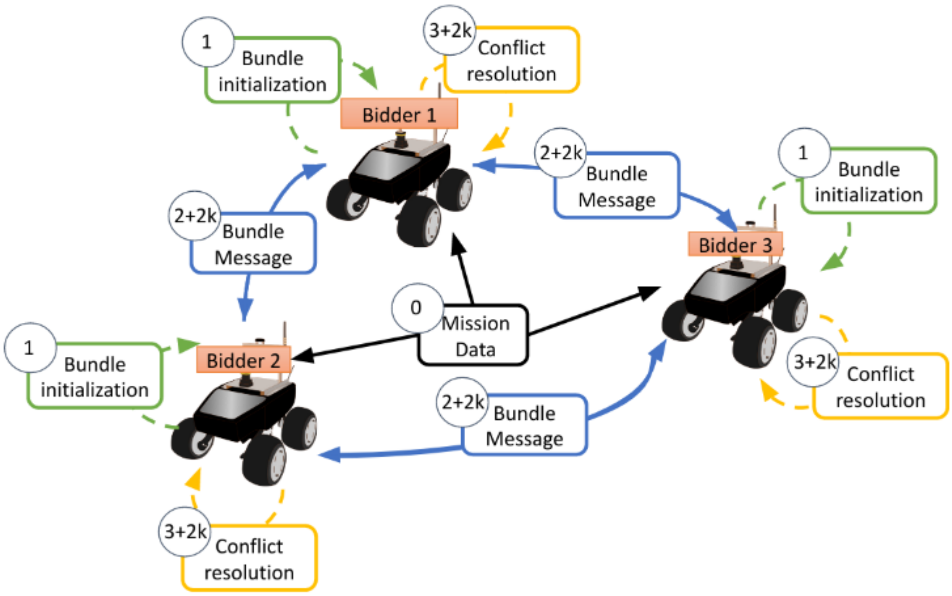

Figure 4: The CBBA auction process [22]. In this figure, bundles refer to bundles (an ordered subset) of tasks. This figure is explained in detail in [6, figure 4]

It is important to note that the two ethos described above each aim to address an extreme of the possible communication network states; continuously fully connected or often fragmented. Those approaches fail to provide an appealing trade-off in a more nuanced application case. For example, while consensus-based methods have proven successful in ensuring robustness in the event of a network failure, a well-designed network might only encounter them sporadically. In such a situation, despite operating in the more optimal network state the majority of the time, the possibility of a network failure forces the solution to adopt an approach robust to such eventualities. In other words, the edge case drives the choice of approach and leads to a sub-optimal process the majority of the time for the sake of reliability.

Based on this observation we propose a novel solution, aimed at seamlessly blending centralised and decentralised architecture to leverage the advantage of both depending on the operational topology (where and when appropriate). This principle, referred to as "**bid intercession**", is explained in section 4.

4 Consensus Based Algorithms with Intercession

Bid intercession (the main contribution of this paper) aims at seamlessly blending centralised and decentralised approaches, with the intent of leveraging the advantages of both ethos where and when relevant, given a network topology state. Intercession can be defined as *the act of intervening on behalf of another party.* In the context of task allocation leveraging market-based approaches, it implies agents evaluating the utility of another, and injecting this information into the decision-making process.

This may be desired in the event that some agents possess certain specific characteristics making them more apt at performing the utility evaluation. These can range from better situational awareness to better technical capabilities or specific responsibilities. The particular scenarios where this dynamic is relevant are addressed in section 6.

The principle of bid intercession is not a consensus mechanism in itself, but rather a re-distribution of responsibilities within a consensus process. It is as such compatible with most existing consensus mechanisms. Bid intercession acts as a "meta"-coordination layer, effectively overriding and re-programming a consensus process based on a given network topology state to follow a different coordination process.

In order to explain the bid intercession mechanics, CBAA[14], one for the consensus-based methods mentioned earlier in subsection 3.3, will be used as the underlying consensus mechanism. CBAA is chosen as it is the simplest variant of the consensus-based algorithms available for task allocation to date. The concept of bid intercession may however be applied to other more advanced algorithms, such as CBBA and does not change in principle to implementation.

CBAA is a task allocation process that operates on a task-by-task basis. In CBAA, each task is auctioned individually and assessed independently of the others. Unlike in CBBA, the valuation of a task in CBAA does not take into consideration the tasks that have already been acquired by an agent. CBAA iterates between two main phases. First, an auction phase, during which each agent estimates bids for the various tasks available (algorithm 1). The second phase is then a consensus process, which is used to converge on a winning bids list (algorithm 3). This logic is follow by every agent in a fleet.

Before diving into the details of the updated algorithm, the MRTA model described in subsection 3.1 must be extended with one more vector. The added vector N_ρ represents the fixed priority level corresponding to each agent in N_u. This vector will serve as the base tie-breaker during the auction process in the event of conflicts. Its role is described in more detail below.

Phase 1: Auction process

The goal of this step is to establish a list of winning bids using an auction process. Each agent considers all the bids placed across the fleet to in turn determine which task it should allocate to itself next. The mechanism as described in [14, Section III-A] of the CBAA paper remains relatively unchanged in principle and simply needs to be extended to enable and support intercession. It functions as follows.

Each agent places bids on the different tasks asynchronously with the rest of the fleet. The auction process is only performed if the agent does not currently have a task assigned to it. The following steps are performed for the task selection:

Step 1 The agent computes a local bid for each task.

Step 2 The local bids are then compared to a list of winning bids to determine the list of valid tasks (for which the local bid outbids the winning bid).

Step 3 The agent then selects the task in the valid tasks list with the highest local bid, updates the corresponding winning bid, and assigns it to itself.

A flow diagram highlighting the process is provided in Figure 5 to help provide an overview of the auction process. Details on the different formulae and vectors are provided below.

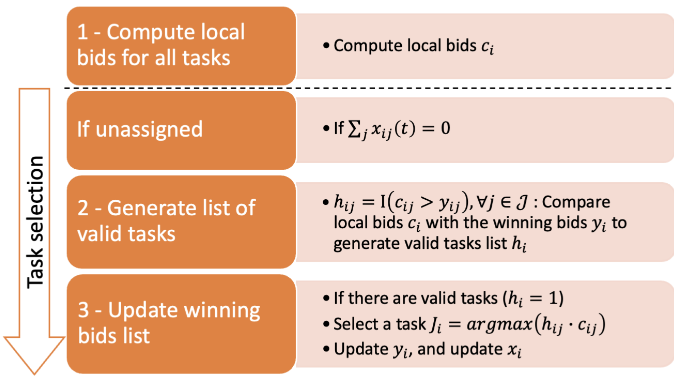

Figure 5: Base auction process flow diagram

$c_{ij} \geq 0$ (computed in Step 1) is defined to be the bid that agent i places for task j. In the original paper, two vectors of length N_t are created (initialised as zero vectors) and maintained by each agent locally throughout the assignment process. The first vector $\mathbf{x}_i$ is the agent's task list, where $x_{ij} = 1$ if agent i has been assigned to task j, and 0 otherwise. The second vector then represents the list of the winning bids $\mathbf{y}_i$. The second vector is assumed to be the most up-to-date estimation of the current highest bids made across all agents thus far (more details on its updating process are provided in the consensus process section below). CBAA being a single-task allocation process, each agent only gets assigned a single task at a time. As such agents only perform the auction process when being currently unassigned ($\sum_j x_{ij}(t) = 0$, with t being time).

The list of valid tasks $\mathbf{h}_i$ can then be generated using the winning bids list (Step 2) as follows:

$$h_{ij} = \mathbb{I}\left(c_{ij} > y_{ij}\right) \quad \forall j \in \mathcal{J} \tag{3}$$

where $\mathbb{I}(\cdot)$ is the indicator function that is unity if the argument is true and zero otherwise. The y_i and x_i matrix are then updated every time the agent wins a task (Step 3), i.e. $\mathbf{h}_i \neq \mathbf{0}$.

The full logic can then be formulated as in Algorithm 1:

Algorithm 1: Auction phase of the CBAA process as described in the original paper

```
SELECT TASK (c_i, x_i(t − 1), y_i(t − 1))
x_i(t) = x_i(t − 1)
y_i(t) = y_i(t − 1)
if Σ_j x_ij(t) = 0 then
    h_ij = 𝕀(c_ij > y_ij(t)), ∀j ∈ 𝒥
    if h_i ≠ 0 then
        J_i = argmax h_ij · c_ij
        x_{i,J_i}(t) = 1
        y_{i,J_i}(t) = c_{i,J_i}
    end
end
```

In order to enable bid intercession, the above-described process needs to be extended as follows. Most of the base algorithm remains the same, with the exception of a few things:

Step 1a During this step, the agent computes local bids for each task, for itself *and each agent it intercedes on behalf of.*

Step 1b A matrix keeping track of the most significant bids placed by across all agents agent (according to agent priorities) for each task is then updated where relevant based on said bids' corresponding priority, the agent's own priority level, and the local bids computed in the previous step. This matrix is referred to as the *current bids* matrix (more below).

The above steps are always performed regardless of whether the agent currently has a task assigned. Then, if the agent is free:

Step 2 The *current* bids matrix (and not local bids) is compared to the winning bids to generate the list of valid tasks

Step 3 Finally, the valid task with the highest bid is selected, and the winning bids are updated.

The resulting extended process is summarised in Figure 6.

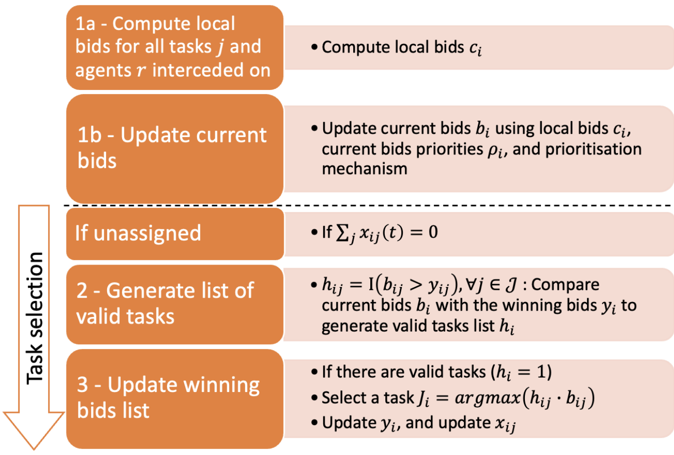

Figure 6: Auction process extended with intercession flow diagram

Accordingly, we first redefine $c_{ijr} \geq 0$ (computed in Step 1a) to be the bid that agent i places for task j, *for agent* r. This results in a matrix c_i of size $N_t \times N_u$. This matrix is initialised as 0 to allow for discerning whether a bid has been calculated or not (it is assumed that the bids resulting from the bid estimation process are always *greater than zero*. Here, $r \neq i$ AND $c_{ijr} \neq 0$ only when agents perform intercessions. The rest of the time the only indexes yielding non-zero values are the ones where $r = i$. The choice of which agent performs intercessions and when is dependent on the choice of architecture and agent hierarchy adopted (more on this in section 6).

The task vectors x_i and y_i remain the same, and the new "current bids" matrix b_i of size $N_t \times N_u$ is defined. It is used to store all *highest priority bids made across the fleet for each task and each agent.* Finally, the ρ_i matrix of size $N_t \times N_u$ is introduced. This matrix represents the *priority levels associated with each bid in* b_i. This priority level is the first element considered during Step 1b of the extended algorithm. It is used to decide whether to account for bid

c_{ijr} or current bid b_{ijr} when updating the b_i matrix, i.e. the priority level $N_\rho(i)$ of agent i is compared to existing priority level ρ_{ijr} associated with current existing bid b_{ijr}. We refer to this process as **prioritisation**.

Note. *When the priority levels are not sufficient to determine the winner ($N_\rho(i) = \rho_{ijr}$), it is necessary to rely on other criteria to break the tie. We refer to those in the algorithms provided as* tie-breakers. *Based on the assumption that more recent bid estimations are likely to be more accurate, we suggest prioritising the bid with the most recent timestamp (resulting in the introduction of an extra N_t vector used to store the timestamp at which a given bid was placed). Other alternatives may however be considered, such as for example a lexicographical heuristic based on the agents' IDs and/or task IDs, as long as they meet two key properties. The following are therefore defined for the prioritisation process to ensure convergence:*

Property 1 (Deterministic and consistent prioritisation)**.** *The prioritisation logic always yields the same results given the same inputs to ensure convergence of the consensus process across all agents in the fleet.*

Property 2 (Globalised priority state)**.** *The prioritisation mechanism used by each agent leverages a prioritisation state globalised across the entire fleet of agents participating.*

Failing to meet either will result in agents obtaining different results in the prioritisation process and in turn locally determining the winning bids inconsistently, preventing the algorithm from converging correctly to a global allocation.

The valid tasks h_i is then generated in Step 2 using the following:

$$h_{ij} = \mathbb{I}\left(b_{ij} > y_{ij}\right) \quad \forall j \in \mathcal{J} \tag{4}$$

It is important to note here that b_{ij} is considered and not c_{ij}. This is necessary to ensure that bid intercession performed by other agents with larger priorities (than that of agent i) are the ones considered in the auction process. Finally, a task is assigned and the winning bids list is updated during Step 3 the same way as it is done in the base algorithm.

The updated algorithm logic is provided in algorithm 2

Algorithm 2: Auction Phase of the CBAA Process extended with Bid Intercession

SELECT TASK $(\mathbf{c}_i, \mathbf{x}_i(t-1), \mathbf{y}_i(t-1), \mathbf{b}_i(t-1), \rho_i(t-1))$
$\mathbf{x}_i(t) = \mathbf{x}_i(t-1)$
$\mathbf{y}_i(t) = \mathbf{y}_i(t-1)$
$\mathbf{b}_i(t) = \mathbf{b}_i(t-1)$
$\rho_i(t) = \rho_i(t-1)$
... (tie-breakers)
# Update b_i
for $\forall j \in$ *rows of* c_i *AND* $\forall r \in$ *columns of* c_i **do**
  if $c_{ijr} \neq 0$ **then**
    $$b_{ijr} = \begin{cases} c_{ijr} & \text{if } N_\rho(i) > \rho_{ijr}(t) \\ b_{ijr}(t) & \text{if } N_\rho(i) < \rho_{ijr}(t) \\ \text{apply other tie-breakers} & \text{if } N_\rho(i) = \rho_{ijr}(t) \end{cases}$$
  end
end
if $\sum_j x_{ij}(t) = 0$ **then**
  # Find valid tasks
  $h_{ij} = \mathbb{I}\left(b_{ij} > y_{ij}\right) \quad \forall j \in \mathcal{J}$
  if $\sum_j h_i > 0$ **then**
    $J_i = \operatorname{argmax} h_{ij} \cdot b_{iji}$
    $x_{i,J_i}(t) = 1$
    $y_{i,J_i}(t) = c_{i,J_i}$
  end
end

Phase 2: Consensus process

The second phase of the CBAA algorithm is a consensus step. The goal of this step is to converge on a single list of winning bids across the agent fleet, which is in turn used to determine the winners and subsequent task allocation.

This in turn allows for a conflict resolution process to be achieved in a network-structure agnostic fashion. Four key steps can therefore be noted:

Step 1 Share local states with neighbours.

Step 2 Receive local states from neighbours.

Step 3 Update local states according to the ones received. from the neighbours.

Step 4 Lose assignment if outbid by neighbours.

The corresponding process is depicted in Figure 7.

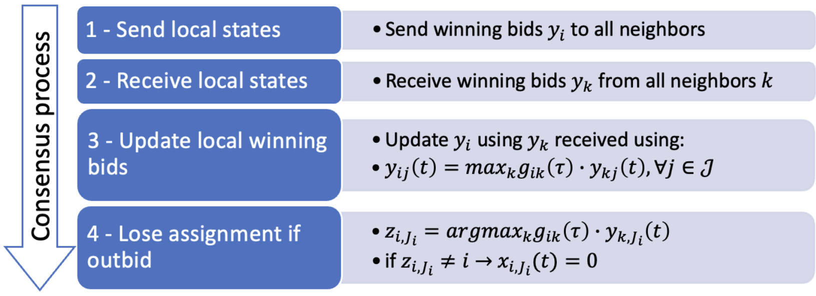

Figure 7: Auction process flow diagram

As described in the section III-B of the original CBAA paper, we start by defining $\mathbb{G}(\tau)$ to be the undirected communication network at time τ with symmetric adjacency matrix $G(\tau)$. The adjacency matrix expresses the existence of links between agents in such a way that $g_{in}(\tau) = 1$ if a link is present between agents i and n, and 0 otherwise. Agents i and n are qualified as neighbours if linked. Additionally, every node has a self-connected edge ($g_{ii}(\tau) = 1 \quad \forall i$ by convention).

The base procedure followed for the consensus is described in algorithm 3, which depicts the agent i's tth iteration when it corresponds to τ in real-time. At each iteration of the second phase, agent i sends (Step 1) and receives (Step 2) the winning bids list y_i to and from each of its neighbours. The consensus process is then performed for each received bid list y_k for all k for which $g_{ik}(\tau) = 1$. In the base procedure, agent i simply replaces y_{ij} values with the largest observed in the y_k obtained from all its neighbours (Step 3). The agent then also loses its assignment if it finds itself outbid by another agent for the task it currently has selected (Step 4). The base algorithm is provided below in algorithm 3

Algorithm 3: Base consensus phase of the CBAA process

SEND $\mathbf{y}_i$ to k with $g_{ik}(\tau) = 1$
RECEIVE $\mathbf{y}_k$ from k with $g_{ik}(\tau) = 1$
UPDATE TASK $\left(\mathbf{g}_i(\tau), \mathbf{y}_{k\in\{k|g_{ik}(\tau)=1\}}(t), J_i\right)$
  $y_{ij}(t) = \max_k g_{ik}(\tau) \cdot y_{kj}(t), \forall j \in \mathcal{J}$
  $z_{i,J_i} = \operatorname{argmax}_k g_{ik}(\tau) \cdot y_{k,J_i}(t)$
if $z_{i,J_i} \neq i$ **then**
  $x_{i,J_i}(t) = 0$
end

Only one modification is necessary here to adapt the consensus phase to support intercession. The updating process must be extended to also merge the received b_k matrix with the local b_i matrix following the prioritisation logic (mentioned previously in the auction phase section). The updated process is shown in Figure 8.

Note that during the updating process, an agent loses its assignment if its local current bids matrix b_i is updated. This is necessary to ensure that in the event of an intercession with a smaller bid, the agent correctly releases the task. The complete updated algorithm is provided below (algorithm 4).

Finally, and similarly to the base CBAA paper and the prioritisation logic, it is assumed that all ties occurring in the determination of either J_i in the auction phase, or z_{i,J_i} in the consensus phase are resolved systematically.

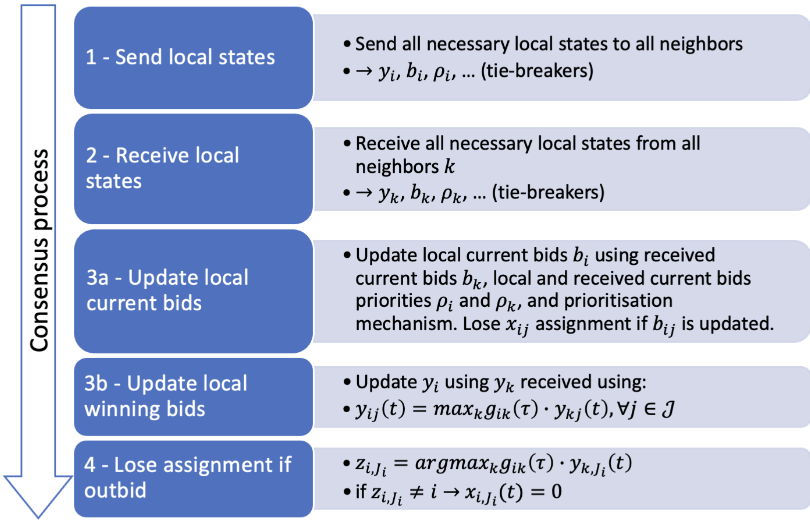

Figure 8: Auction process flow diagram extended with intercession

Algorithm 4: Base consensus phase of the CBAA process

SEND $\mathbf{y}_i$, $\mathbf{b}_i$, ρ_i, ... (tie-breakers) to k with $g_{ik}(\tau) = 1$
RECEIVE $\mathbf{y}_k$, $\mathbf{b}_k$, ρ_k, ... (tie-breakers) from k with $g_{ik}(\tau) = 1$
UPDATE TASK $\left(\mathbf{g}_i(\tau), \mathbf{y}_{k\in\{k|g_{ik}(\tau)=1\}}(t), \rho_{k\in\{k|g_{ik}(\tau)=1\}}(t), J_i\right)$
for $\forall j \in$ *rows of* b_k *AND* $\forall r \in$ *columns of* b_k **do**
  if $b_{kjr} \neq 0$ **then**
$$b_{ijr} = \begin{cases} b_{kjr} \text{ and } x_{ij}(t) = 0 & \text{if } \rho_{kjr}(t) > \rho_{ijr}(t) \\ b_{ijr} & \text{if } \rho_{kjr}(t) < \rho_{ijr}(t) \\ \text{apply other tie-breakers,} & \text{if } \rho_{kjr}(t) = \rho_{ijr}(t) \end{cases}$$
  end
end
$y_{ij}(t) = \max_k g_{ik}(\tau) \cdot y_{kj}(t), \forall j \in \mathcal{J}$
$z_{i,J_i} = \text{argmax}_k\, g_{ik}(\tau) \cdot y_{k,J_i}(t)$
if $z_{i,J_i} \neq i$ **then**
  $x_{i,J_i}(t) = 0$
end

5 Method Analysis

This section focuses on studying the impact of introducing intercession in consensus-base allocation algorithms. The analysis concentrates on two key aspects: convergence (subsection 5.1) and performance (subsection 5.2). Those are the sole fundamental properties necessary to be understood in order to establish the fundamental viability and potential of the principle of intercession in the context of such applications. Convergence evaluation delineates the ability of the algorithm to reach a stable solution, while performance scrutiny aims to demonstrate the potential impact of the method on the resulting solutions' quality. Where possible, numerical results are provided to assist in the verification of claims and analysis. The details of the experimental setup and configuration corresponding to each test can be found in section 8 and the code can be found on github[2].

5.1 Convergence

The algorithm's convergence refers to its ability to produce a conflict-free assignment in a finite amount of time. Two key aspects are considered here: the ability to converge, referred to as *convergence termination*, and the *convergence iteration complexity* i.e. the number of logical steps necessary to converge to a solution. Additionally, it is also important to define the *time complexity* as the time required to converge to a solution. In order to evaluate those properties, we start by demonstrating that it is possible to inherit the convergence guarantees of the underlying decentralised algorithm (property 3). Then, to verify the reorganisation of responsibilities capability enabled by intercession (property 4), the most extreme case study is considered; a fully centralised configuration, where a single agent intercedes on behalf of all others (the One-to-Many architecture, see section 6). Verifying the above will ensure

[2] https://github.com/vguillet/AE5310_Master_Thesis

that a clear understanding of the algorithm's operational properties is established for both extremes of the coordination spectrum (fully decentralised, and fully centralised).

Before evaluating the hypotheses proposed for this analysis, the following must first be defined:

Lemma 1 (Unambiguous source of truth)**.** *The prioritisation process yields consistent results across all agents when performed on the basis of the same information.*

Proof. In the above-defined problem and algorithm, the fixed nature of the agents' priorities (property 2) combined with the deterministic and constant nature of the prioritisation mechanism (property 1 defined in section 4) allows for asserting the deterministic and unambiguous nature of the source of truth at all times. □

Property 3. *The process of bid intercession does not have any impact on the convergence of the underlying distributed method of a given auction, given that the source of truth is unambiguously defined at all times. More specifically, the following are not impacted:*

- *Convergence termination*
- *Convergence iteration complexity*

Proof. As specified in [14, section V-D]: *"[...] whatever knowledge each agent scoring scheme is based on, the only needed information for resolving conflicts among agents are the winning bid list, winning agent list, and the time stamp. If these three pieces of information are communicated error-free,* ***the conflict resolution process of (CBBA) is <u>insensitive</u> to the details of each agent's scoring scheme."***

The specific algorithm analysed by this statement (CBBA) is placed in brackets as the statement also applies to the CBAA algorithm [14, section III].

Bid intercession solely introduces additional logic and conditions adjusting the source of information in the process of winning bid determination and exchange of information (intercession requiring slightly more already existing information to be passed around). These modifications in turn allow for intelligently managing different scoring schemes present within a single network, and "overwriting" scoring results before performing the auction process (which itself remains unchanged). Put simply, if an agent x places a bid on behalf of agent y with a larger relative priority, it amounts to the exact same as if agent y had placed the same bid itself in the CBAA framework. The detail of the scoring process logic and/or the data used in it is therefore abstracted away in the bid value. Intercession thus fundamentally only effectively modifies the agents' scoring scheme. Accordingly, the core process of reaching consensus and resolving conflicts, as defined by the underlying algorithm, remains unaltered and unaffected.

Consequently, the above provided in conjunction with lemma 1 makes it possible to infer that the introduction of bid intercession has no impact on the convergence properties of the underlying algorithm it uses as a baseline, thereby verifying property 3.

To further verify this property, a number of simulations were performed with the goal of comparing the convergence properties with and without the introduction of bid intercession. The results confirmed that the extra logic did not impact the underlying algorithm convergence termination in any way as expected. Details of the simulation configuration and results can be found in subsection 8.3.

Additionally, in the context of CBAA enhanced with bid intercession (and given lemma 1) the following statement can therefore also be made:

"Whatever knowledge each agent scoring scheme is based on, the only needed information for resolving conflicts among agents are the current bids matrix, current bids priority matrix, and any other data points necessary for the prioritisation process. If these pieces of information are communicated error-free, the conflict resolution process of (CBAA) ***with intercession*** *is insensitive to the details of each agent's scoring scheme.*"

It is important to note here that while intercession does not impact the consensus termination and iterative complexity, the same cannot be said about the *temporal complexity* of the algorithm. The introduction of additional bid estimation algorithms with more or less complex algorithms will have an impact on computation time, which may impact the time at which the final consensus is reached. □

Property 4. *In a connected static network and a heterogeneous fleet with a single higher priority agent bidding for all others based on a centralised allocation method, the method produces the same results as the centralised method (used by the higher priority agent).*

Proof. The lower bound of the coordination centralisation range (fully decentralised) is guaranteed by the underlying algorithm, CBAA and CBBA having been designed specifically for this case. To verify the other extreme (fully centralised), the following conditions are considered:

1. All agents reach the same conclusions on the basis of the same prioritisation information (lemma 1)
2. The same prioritisation information is correctly provided to each of them (which can be affirmed as a direct result of the connected network assumption)

Additionally, the following network configuration is established:

3. A single agent with the highest relative priority emits bids for all other agents reflecting an allocation plan centrally determined (a design choice)

Combining all of the above makes it possible to induct that in such a context, the final allocation will correspond to the one determined centrally by the high-priority agent, thereby confirming property 4. This was yet again verified experimentally, and the results do confirm that the network consistently converges around the centrally determined solution (more details on the experimental setup and results can be found in subsection 8.4). □

5.2 Allocation Performance and Outcome Manipulation

The allocation performance refers to how effective the final allocation is at maximising/minimising a specific metric (see subsection 3.1). Bid intercession offers the possibility of further tailoring the allocation logic to a specific scenario and network configuration. This, in turn, aims to further enable the utilisation of highly heterogeneous fleets and efficiently leverage various resources distributed unevenly throughout the network. Accordingly, the actual final allocation quality greatly depends on the scenario of application, choice of architecture, and agent hierarchy implemented in the system. It is as such difficult to establish a clear magnitude for the potential gains in performance enabled by bid intercession. It is however possible to demonstrate that the allocation performance can be manipulated.

To do so, the following demonstration is proposed. Two categories of agents are established:

- **Base agents:** Those agents operate using the simplest (relatively speaking) bidding mechanism to estimate paths and their own utility. Those agents do not intercede on behalf of others and solely focus on their own role in the process. They hold the responsibility of completing the tasks. A real-world analogy would be small drones with limited sensors and onboard processing power.
- **Interceding agent:** Those agents have access to a more complex bidding mechanism. They are capable of leveraging more information in the estimation of their utility. Those agents are not mobile, and as such cannot take on tasks. They are solely present to assist in the coordination process and possess a higher priority level than the base agents. The interceding agent will be bidding on behalf of all other agents, resulting yet again in a One-to-Many architecture. A real-world analogy of such agents would be a control centre, or a van brought on-site. Those would have more room for extra computational power, better communication, and access to more data. *Note: In this experiment, a single interceding agent is used.*

The following scenario is then established. Agents are responsible for allocating among themselves and taking on tasks emitted at unknown and varying intervals. The agents will operate in a connected grid environment, where each existing edge e_{ij} connecting two nodes n_i and n_j represents a step with an associated cost e_{ij_c} (with a value in the $0-100$ range). The degree of interconnectivity present within the grid can vary from one scenario to another. The tasks are locations in the grids that must be visited. An example of such an environment can be seen in Figure 9.

The quality of the allocation will be evaluated on the basis of the cumulated step cost taken by the whole fleet to complete all tasks. The base agents' bid will be estimated as the inverse of the Manhattan path $\text{MP}_{n_a \to n_b}$ from the current agent i's location a to a target location b, with each step having a cost of one (the base agents being effectively blind to the environment step costs). The bid formula can be seen in Equation 5.

$$\text{bid}_{iji} = \frac{1}{\sum_{e_{ij} \in \text{MP}_{n_a \to n_b}} 1} \tag{5}$$

The interceding agent's bid will be estimated using the inverse of the weighted Manhattan path $\text{WMP}_{n_a \to n_b}$ between an agent r's current location a and a target location b, allowing for bids which take into account the environment's step costs e_{ij_c}. The corresponding formula is shown in Equation 6

$$\text{bid}_{ijr} = \frac{1}{\sum_{e_{ij} \in \text{WMP}_{n_a \to n_b}} e_{ij_c}} \tag{6}$$

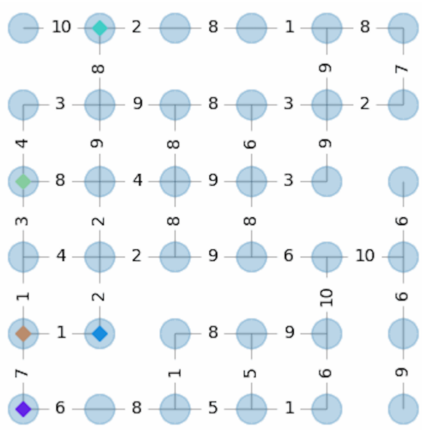

Figure 9: Example environment used in the experiments. This environment contains 36 nodes and is 80% connected. The coloured diamonds visible on it represent the different agents present.

For this scenario, an extra set of data points will be provided with each bid (akin to the bid priority). That is, the path $\mathrm{MP}_{n_a \to n_b}$ or $\mathrm{WMP}_{n_a \to n_b}$ corresponding to each respective bid. This is necessary to ensure that an agent knows which path to follow upon winning a task with a specific auction. To keep the bids and paths up to date, the allocation process was set up in such a way that each agent recomputed their bids at every change of state.

A Monte-Carlo approach is adopted to test a large array of configurations of the above-described setup, with each scenario here being tested in pairs; one run only contains base agents and the other contains the base agents and a single intercession agent. The task announcement schedule and environment are kept constant within each pair, resulting in the only difference between the two being the fleet configuration and hierarchy. The exact details of the experimental setup and simulation configuration can be found in section 8. For each pair, the final cumulated cost of each respective run is computed upon completion, and the difference is determined. The results of 1070 runs can be observed in Figure 10 and Figure 11.

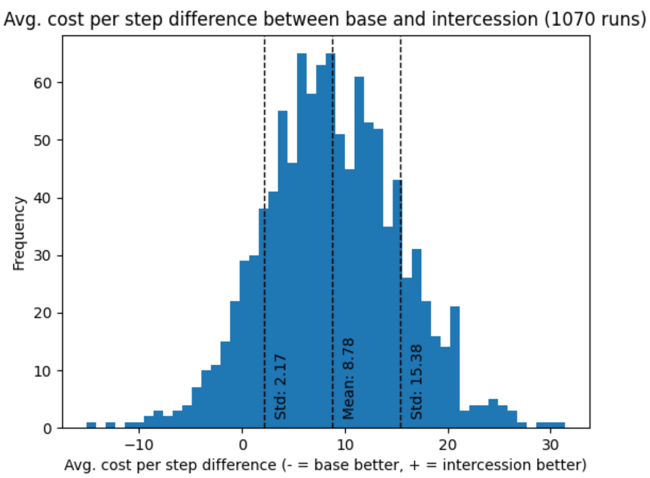

Figure 10: Distribution of results for each run pair. It can be observed that the average cost per step is lower with intercession, highlighting the positive impact of intercession on the scenario

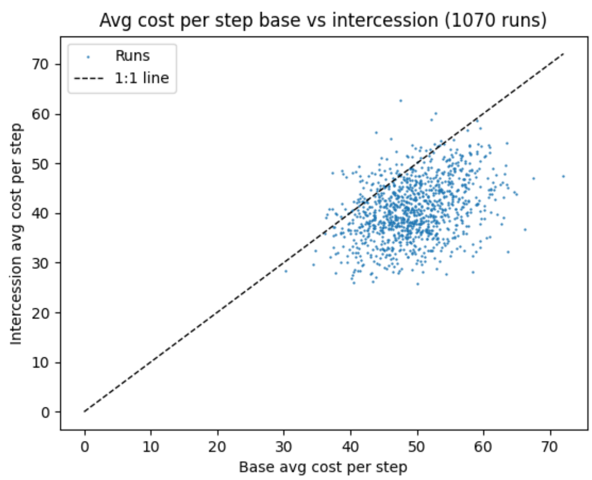

Figure 11: Plot of the average cost of a step for each run for the base run and intercession run in a pair. The large majority of points are located below the 1:1 line, indicating that steps are on average more costly for the base runs compared to their intercession equivalent

As can be seen in the simulation results, the process of intercession enables us to successfully manipulate the outcome of a task allocation. Note again that the magnitude of the bias introduced in the allocation process is dependent on the specific architecture implemented in the agent fleet, and the potential corresponding gains present in the problem itself.

It may be noted that the results demonstrate that such intercession will occasionally fail to improve the quality of the final allocation. This is a direct result of the temporal myopia nature of the allocation process. In other words, a more optimal solution at an instant t with respect to a specific goal does not necessarily ensure a more optimal final result.

In this context, taking a different path might lead to an agent being less favourably positioned relatively when a new task appears. This can however be improved through the implementation of more advanced allocation methods (for example CBBA instead of CBAA), and/or more advanced bidding mechanisms (accounting for environment "coverage" in anticipation of new tasks for example).

6 Applications

This section mainly focuses on presenting possible implementations and variants of consensus-based task allocation with intercession in the context of the SAR case study.

The principle of intercession is before anything a dynamic reorganisation of responsibilities within a single agent network. Through the use of priorities, agents are able to override other agents' evaluation of their respective utility, and in turn, take control of part (or all) of the consensus process. This opens up the possibility of adjusting the level of centralisation desired and the structure of the allocation process while always ensuring a resilient system in the event of agents or communication failure. An initial concept must however first be presented before detailing the various allocation configurations possible:

Restricted Auction: Intercession opens up the possibility of controlling the decision-making capability of different agents while still retaining a unified algorithmic baseline. Through the introduction of additional constraints (such as flags, or minimum-priority level requirements), and some slight modification to the auction mechanism to handle the edge case of no bids being present at times (resulting in no one winning the auction), it is possible to limit auctions to a specific subset of agents. A restricted auction is therefore an auction with a set of requirements and criteria for allowing participation. Although already feasible in the original consensus algorithms, this concept significantly gains in potential with the introduction of intercession principles. Examples of applications highlighting this are provided in the analysis of the different possible agent architectures below.

The design options related to which (and how many) agents intercede on behalf of which, in which context, and for what reasons are extremely varied. A number of notable fleet organisation schemes can however be isolated. Those are:

6.1 *One-to-Many:* A single agent interceding on behalf of all other agents.

6.2 *Many-to-Many:* Multiple agents interceding on behalf of multiple/all other agents

6.3 *Many-to-One:* Multiple agents interceding on behalf of a single agent.

6.4 *One-to-One:* A single agent interceding on behalf of a single agent.

6.5 *Combined Architecture:* Combining all of the above.

Those are described in more detail below and contextualised with respect to the SAR problem.

6.1 One-to-Many: A single agent interceding on behalf of all other agents

This configuration results in the centralisation of the decision-making process and is key here to enable the use of centralised algorithms within a decentralised method. Given a set of tasks, a central agent performs a local allocation of all the tasks across all the agents using *any arbitrary centralised allocation algorithm*. The allocation here is performed on the basis of data observed by the central agent, and/or possibly local data relayed back by the agents on the field. The resulting allocation is in turn "translated" into the corresponding bids, with larger bids for each winning agent-task pair matching the allocation plan, and small ones for all others. The bids are then emitted with a larger priority level, ensuring that the central agent plan is prioritised over all local estimations. The central plan is then propagated to all agents using the auction/consensus phases, and results in the desired allocation as defined by the central agent. A graphical representation of the architecture can be found in Figure 12.

It is important to note that here the effectiveness of the centralised method is still dependent on the topology of the communication network. The plan as defined by the central agent will only be respected by the agents that the instructions are able to reach. If a sub-group of agents finds itself cut off from the main group for an extended period of time, the allocation process will continue there according to the local bid estimation performed by the agents until reconnected with the network. At this point, a re-planning phase by the central agent might be necessary to account for the fact that some of the agents were not able to perform as expected, resulting in a slightly altered allocation. Going back the the SAR case study, this principle could be implemented through the use of a central planner, located in a headquarters or simply in the cloud (assuming the means to connect to this cloud infrastructure and relay the instructions is available on-site and is still functional).

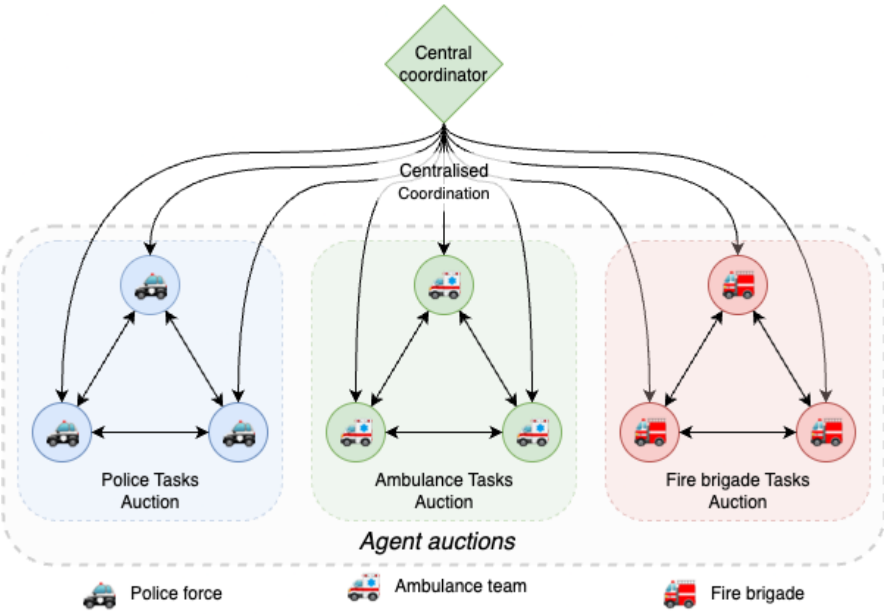

Figure 12: The One-to-Many architecture

6.2 Many-to-Many: Multiple agents interceding on behalf of multiple/all other agents

This configuration might prove useful in multiple scenarios and for different reasons. A first example would be to increase redundancies in the previously mentioned centralisation process. Having multiple agents performing a centralised allocation with different priority levels ensures that if the main planner fails or becomes unreachable, the others may act as backup sources of coordination. Note that if all of those in turn fail, the system naturally falls back to local bid estimations, ensuring a resilient system at all times. This is also an example of the Many-to-One architecture (detailed further down) when looking at this configuration from the perspective of a single agent being the target of the intercessions. A graphical representation of the architecture can be found in Figure 13.

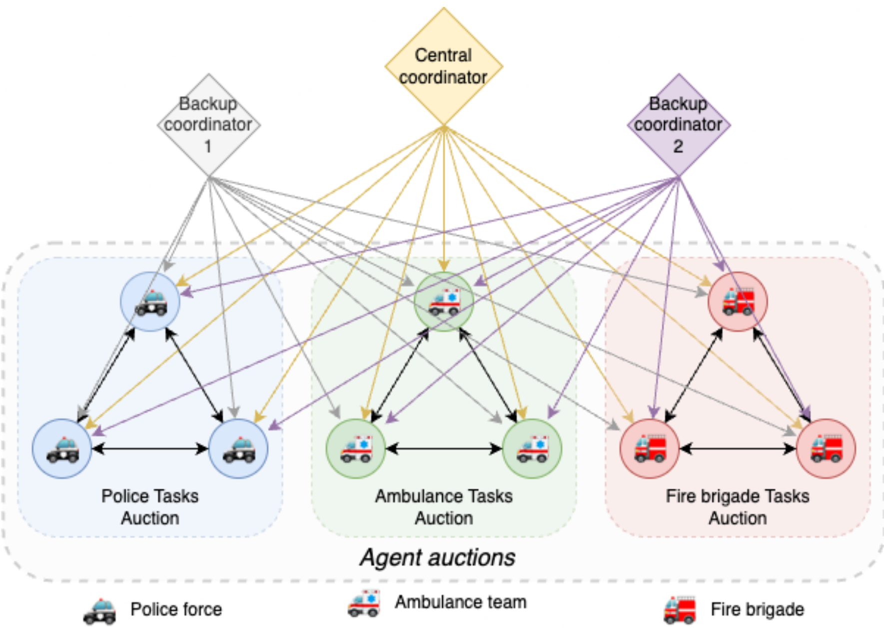

Figure 13: The Many-to-Many architecture can be leveraged to build redundancy in the centralisation process

Another such application is tiered coordination. Some agents could operate in units, in which a single agent takes on the responsibility of "team leader". The team leader would be in charge of locally centralising the coordination of "agent tasks" (all but the coordination tasks listed in section 2). Tying back to the SAR case study, the location agents (ambulance centre, fire station, and police station) could take on the responsibility of coordinating their respective mobile agents centrally. The location agents would then be responsible for communicating among themselves in order to not only share information, but also establish the most efficient strategy to address the situation, and in turn, publish an allocation to their respective agents on site. A diagram representing this configuration can be found in Figure 14.

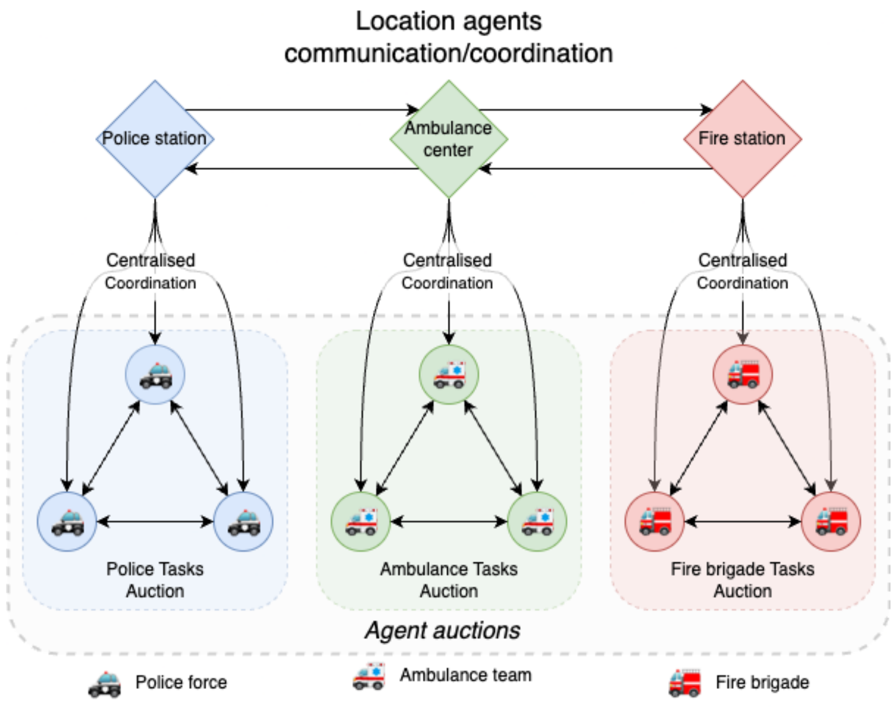

Figure 14: The Many-to-Many architecture can be leveraged to create multiple central coordinators which coordinate themselves/share information to reach their respective allocation

This architecture can be integrated with the previously described restricted auction principle, allowing not only for the consolidation of all consensus sub-networks into a unified whole but also for the seamless design and integration of command chains directly within the allocation process.

Going back to the SAR example, we will for this example define the team leader agents to be specific mobile agents (police forces, fire brigades, and/or ambulance teams). *Note: The location agents could be used as team leaders again, but it is chosen as such here for the sake of showcasing yet another possible configuration where heterogeneous teams are used (multiple different classes of agents).* The first step to introducing an additional layer of coordination is to insert a team-leader-only (restricted) auction in the consensus process. The purpose of this auction will be to allocate groups of tasks among the different participating teams (this auction will therefore be referred to as a *task groups auction*). Tasks may be grouped based on a variety of criteria and conditions, and team leaders will assess bids for these task groups to reflect the team's overall suitability for the specific defining features of each task group. For instance, when dealing with task groups situated within a particular area, team leaders may calculate their bids by considering their team's centroid position in relation to one of the tasks within the group being auctioned off. The allocation of the task groups effectively serves the sole purpose of *helping ponder the eventual bid estimation performed by the members of each respective team*. In other words, the bid estimation performed by each agent is ponderated according to the outcome of the task group auction. The task group auction is therefore used as an intermediary planning step in order to in turn determine a final agent task allocation, and effectively can be seen as embedding decision trees directly in the allocation process. A visual representation of the architecture at this stage is shown in Figure 15.

The second step is to introduce centralising roles in the network. As explained in the first example of tiered coordination (Figure 14, we can give certain agents the responsibility of coordinating their respective groups. Based on the outcome of the task groups auction, the team leader is made responsible to intercede on behalf of the members of their respective teams and determine the allocation of the tasks in the task group, while placing a very low bid for all others. A diagram representation of this setup is provided in Figure 16.

Such architecture presents a number of significant advantages. Conditioning the bid estimation of the team leader agents (for their respective team's agents) on the basis of the outcome of the team-leader auction effectively amounts to constructing and embedding command chains directly in the decision-making process, further allowing for tailoring the allocation logic. This may be further combined with centralised coordination as, by interceding in the team leaders auction, a general coordinator may effectively pilot the teams with broad goals (task groups), while leaving to the

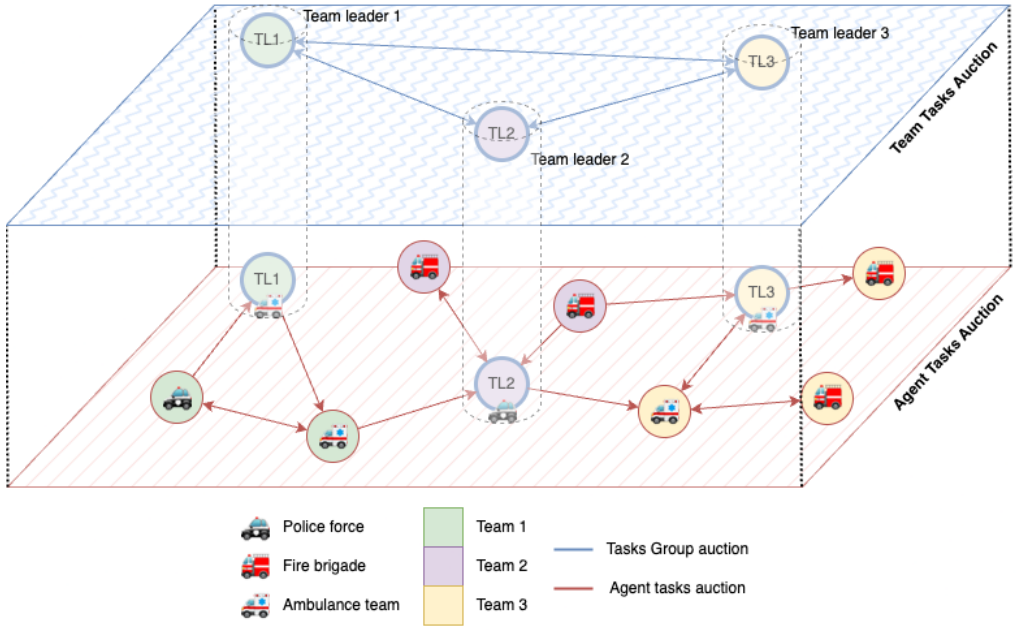

Figure 15: Restricted auctions enable nesting multiple auction processes within a single unified consensus network, allowing in turn for an allocation process to be broken down into multiple interlinked auctions

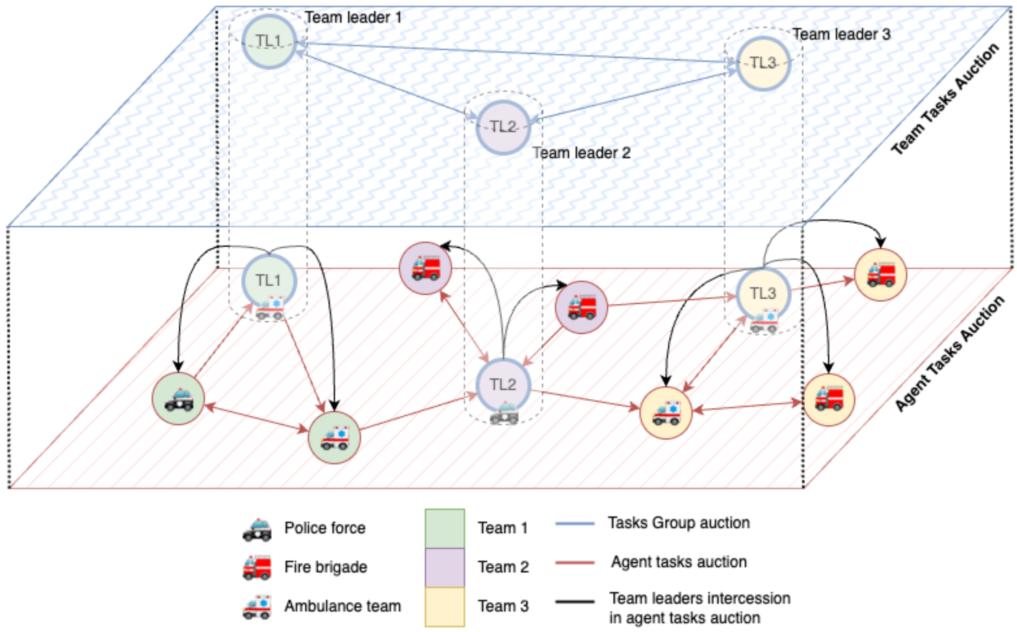

Figure 16: Combining Many-to-Many architectures, Many-to-One, and restricted auctions enable for embedding decision-trees directly in the allocation mechanism

teams on the field the responsibility of allocating the specific tasks among themselves. This is particularly useful in scenarios where a general strategy is devised, and where the team themselves need to retain operational autonomy in the allocation of the tasks requiring on-the-field information. Note that, while it may be possible to introduce multiple interlinked auction steps to determine an allocation using consensus mechanisms without intercession, achieving centralised coordination around team leaders is impossible without it. Additionally, as mentioned earlier, the use of restricted auctions allows for performing all of the different auctions within a single unified consensus framework. All agents act as relay nodes in the network, regardless of their participation in the different auctions.

The architecture described so far however is not yet fully robust to network failure. In the event of one of the team leaders being completely cut off from the network, the lack of representation of the team in the task groups auction leads to the team being completely left out of the allocation (as all task groups would be assigned to the other teams). As such, the team leaders effectively become single points of failure for their respective teams in this structure. This can however be addressed using intercession through the introduction of an additional architecture: Many-to-One.

More on this is provided lower down.

While the example provided here only introduces a single coordination layer, there is theoretically no limit to the choice of structure and number of coordination steps which could be included in a coordination architecture. It can however be noted that a more complex coordination structure will likely lead to longer convergence times, as a more complex and dynamic consensus is consequently devised. This however remains highly dependent on the structure and nature of the decision process devised. The impact of the design of consensus architecture on convergence time is however not a focus of this paper, and as such was not investigated. Nevertheless, such configurations highlight the distribution of responsibilities and flexibility in the coordination/centralisation structure enabled by the principle of intercession.

6.3 Many-to-One: Multiple-agents interceding on behalf of a single agent

This configuration results in the auction being effectively conducted through the prioritisation process and leads to only accounting for the priority levels instead of the bids for a given task. This becomes important in contexts such as the tiered task allocations described above in the Many-to-Many configurations. Going back to the second tiered auction example (Figure 16), it is necessary to ensure that a representative of each team is always able to take part in coordination auctions in order to ensure that the process remains robust to communication failure. This is critical in the event of a complete loss of contact with the team leader agent. *Note: It is important to note here that while intercession enables agents to bid on behalf of other existing agents, it also makes it possible to bid on behalf of fully virtual ones. In other words,* ***an agent may be solely present in the fleet as a reference in the auction ledger*** *and used as a pointer for representing an abstract entity.* In our tiered auction example, the task group auction may effectively be conducted between virtual entities (a.k.a. auction ledger entries) through intercession, each representing a team. We will call them *team agents*. The same original architecture may thus be reproduced by adjusting the role of the team leaders to be interceding for their respective team agents, which acts as a proxy representation in this specific auction process. At this point, leveraging the Many-to-One architecture introduces the resilience missing from the original example. Instead of the team leader agent being the team representative, we can make some/all agents in a team intercede on behalf of their respective virtual team agent, with a defined priority hierarchy. This ensures that the aptest agent always retains priority while allowing other agents to progressively "take on" the role of team representative, following a chain of command in the event of the various team members encountering catastrophic failure. Note here that not all (if any) agents in a team need to possess the capability to centralise the allocation for their respective (or even whole) team. The main difference with respect to the architecture without the One-to-Many process is that the team is always represented in the team leader auction, and as such will still get allocated task groups (irrespective of the coordination process specifics present within each respective team). In the event that no other agent but the primary team leader has the capability of centralising coordination within a team and this very agent goes missing, the tasks in the task groups will fall back to being allocated using the team agents' local bid logic. A graphical representation of such a scenario is provided in Figure 17.

6.4 One-to-One: A single agent interceding on behalf of a single agent

A real-life example of such a scenario would be a human operator providing a specific instruction: "ambulance team A must go to location (x, y)". Through the process of intercession, injecting such instruction into the decision-making process proves to be trivial. In such a context, the human operator can be considered as an agent in the network. The instruction is converted into a new GOTO (x, y) task, and a bid with the maximum priority and bid value possible is emitted (one could go a step further and introduce further multi-level bid logic such as flags in the prioritisation mechanism to ensure that the human-generated instructions are always prioritised). The new task and bid then propagates throughout the network and results in the desired allocation. The reverse is also possible, as an extremely small bid injected can ensure that an agent does not obtain a specific task. Such concepts can furthermore be extended with the principle of a restricted auction to ensure that certain decisions/allocations remain the sole responsibility of a specific actor (such as a human). In other words, certain tasks may only be allocated and undertaken through human intervention for example. This may prove critical in situations where more complex ethical considerations or nuanced evaluation of a situation are required. Note that here any task may still be discovered and announced by any agent. Such scenarios highlight the fine level of external control enabled in the allocation process.

6.5 Combined Architecture: Combining all of the above-mentioned architectures

Having exposed the key allocation architectures enabled by intercession, it is now important to highlight that all of them are compatible with each other. Figure 17 serves as a compelling example showcasing both the compatibility and strength of the combined architectures. To further expand on this, we can introduce a central coordinator who will

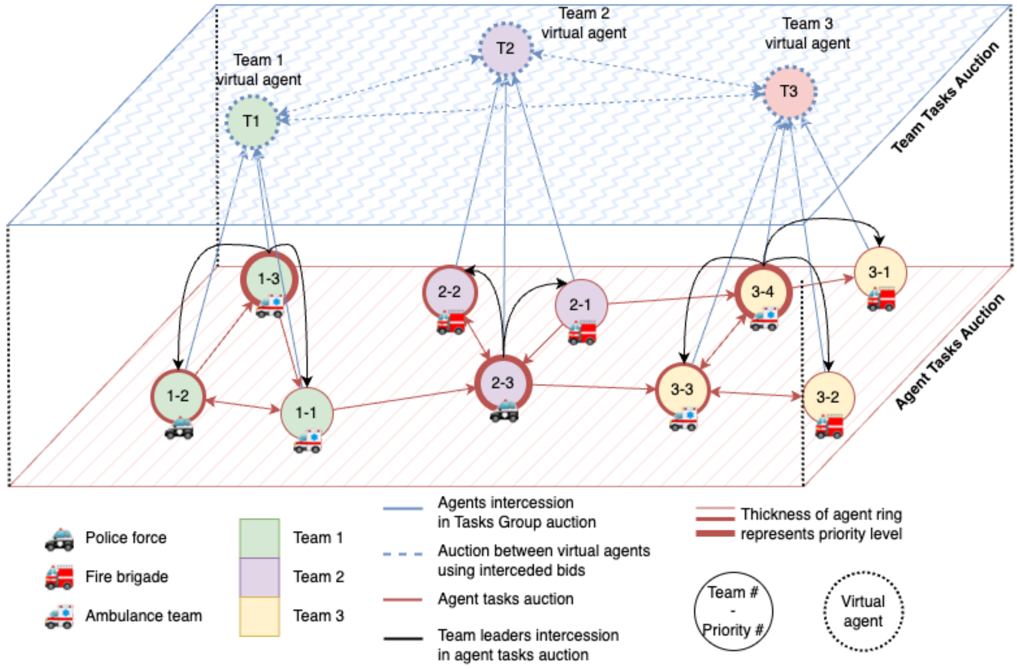

Figure 17: Combining Many-to-Many architectures, Many-to-One, One-to-Many, and restricted auctions enable for the **robust** embedding decision-trees directly in the allocation mechanism

serve the role of coordinating the entire fleet through the emission of general instructions in the form of task groups (One-to-Many). Additionally, we will assume that a control centre managed by humans (referred to as supervisors) monitors the entire process, and occasionally intervenes to correct the allocation by interceding on behalf of the different agents when necessary (One-to-One). It is also assumed that humans always possess the highest priority across the whole system. A final representation of the entire architecture can be seen in Figure 18.

Applying the above-described architecture to the SAR case study allows for observing a number of key advantages:

- The SAR coordination process can be designed to follow the inherent structure of the teams operating on the field. This, in turn, allows for enabling realistic and effective coordination measures, while still staying compatible with existing hierarchies and operational principles (possibly reducing retraining necessary for example).
- The coordination process implemented is robust and resilient enough to withstand the unreliable and dynamic context of operation inherent to the SAR case study.
- The coordination process is flexible and controllable. It is simple to intervene in the process where and when necessary to manipulate both bids and task announcements.
- The allocation process allows for the seamless collaboration and interaction of humans and machines.
- The allocation process is programmable and efficient. It becomes possible to leverage various algorithms and distributed resources in order to seek an effective response given the operational context and real-time constraints.

7 Discussion

This section serves the purpose of discussing the various aspects of the algorithms presented until now, and pointing out the key areas of interest.

Before all, it it is important to highlight that the choice of structure for the algorithm was conceived to ensure that the consensus was capable of operating in a near-identical fashion to the original algorithms. Additionally, the proposed algorithm was specifically selected to present the simplest possible functional implementation of the principle of intercession. This consequently leaves plenty of room for improvements at various levels, which may in turn significantly improve the properties of such allocation algorithms. Those are not the subject of this study and are therefore left for future work. The introduction of intercession however brings forth seven critical areas for discussion:

7.1 *Hybridised Allocation Mechanisms:* Combining multiple allocation mechanisms

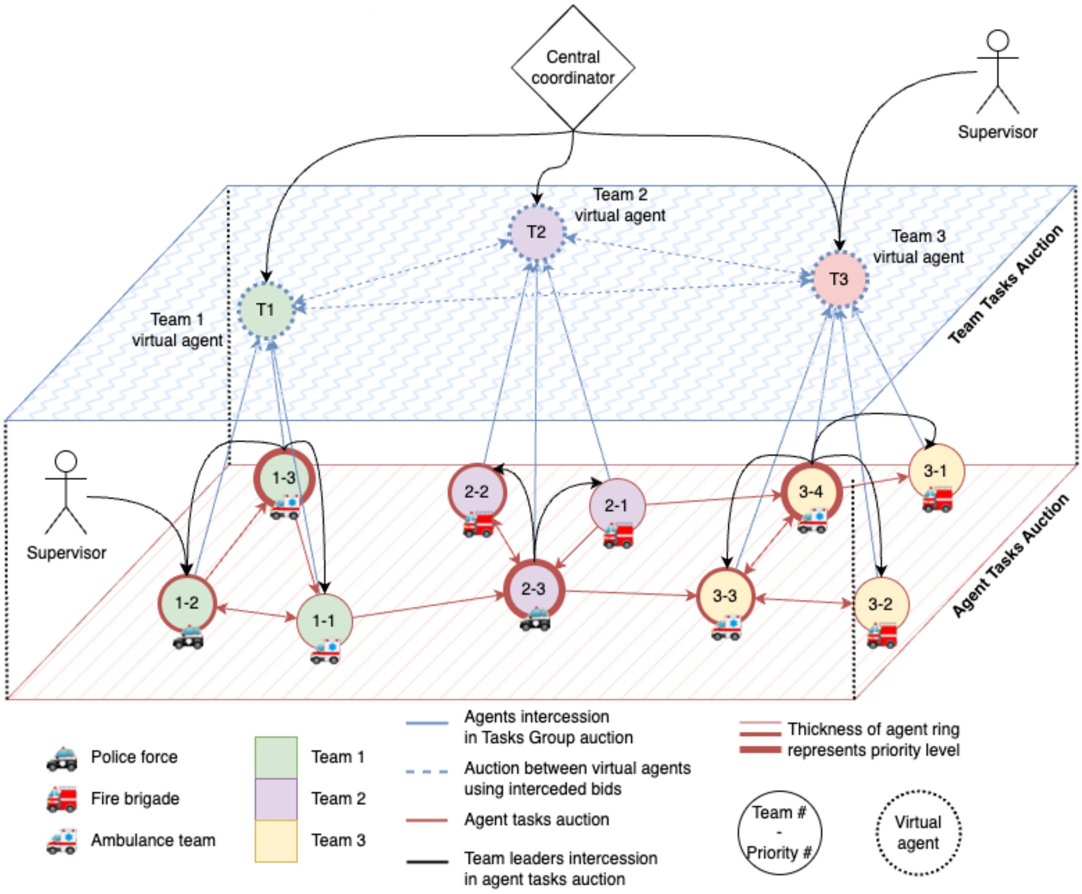

Figure 18: Combining One-to-Many, Many-to-Many architectures, Many-to-One, One-to-One, and restricted auctions enable for a highly complex, robust, and flexible allocation mechanism

7.2 *Tiered Task Allocation:* Embedding hierarchies and decision trees in the decision-making process

7.3 *External Control:* Enabling external control of the allocation-process

7.4 *Integrated Responsibilities:* Embedding accountability at the allocation level

7.5 *Communication Optimisation:* The impact of the modification on the communication loads

7.6 *Advanced Prioritisation Mechanism(s):* The importance of the prioritisation mechanism

7.7 *Failure modes:* The key failure modes possible for such systems

These will be discussed in greater detail below, serving as a starting point for further exploration.

7.1 Hybridised Allocation Mechanisms

Efficiently leveraging available resources is critical to ensuring the solid performance of an agent network when tackling a specific scenario. Agents must be able to not only effectively take advantage of the available resources, but also adapt to changes in real time to make the best of any given situation. The operational conditions must also be considered, such as variations in the fleet structure, agent count, and communication network integrity. Combining consensus-based mechanisms with intercession helps here by enabling the integration of multiple consensus, auction, and bidding mechanisms seamlessly (more in section 6), allowing for the implementation of adaptive coordination mechanisms. The choice of which allocation mechanisms and bidding logic to use in which context, and with which condition remains to be explored. It may however be supposed that correctly choosing the different elements making up an allocation strategy may possibly significantly boost flexibility and performance within the correct context of application. Additionally, implementing the intercession principle in other consensus-based task allocation algorithm variants such as CBBA should also be investigated.

7.2 Tiered Task Allocation

Intercession combined with other concepts such as restricted auctions (see section 6 below) enables the implementation of complex decision-making hierarchies. Those in turn may be structured to follow complex decision trees and flows, all the while remaining distributed in operations. This may prove particularly valuable in large organisations requiring more structured coordination to operate efficiently. This is yet again another area of research which may prove fruitful, not only in terms of performance but also with regard to the features and flows which may be enabled.

7.3 External Control

The principle of intercession opens up the possibility of directly intervening in the task allocation process through the emission of high-priority bids. The One-to-One architecture (see section 6) best highlights this concept. Such intercession enables an arbitrary participant (e.g. a human operator) to directly and punctually influence the allocation process without impacting the underlying mechanism.

This in turn paves the way for efficient and effective human-machine interactions in such systems and introduces the ability to incorporate human expertise, preferences, and real-time insights into the task allocation process. This human-centric approach may therefore ensure that the multi-agent fleet operates with the precision and adaptability that only human judgment can provide at times.

7.4 Integrated Responsibilities

The external control enabled by intercession allows for the seamless integration of responsibilities into the decision-making process. In other words, decision-making authority can be specifically assigned to different stakeholders in the system. This in turn can ensure that the various decisions are taken by the correct participants best able to contribute in a controlled fashion to the process, enhancing efficiency and *accountability*.

This new method consequently facilitates aligning systems with current regulations for autonomous systems in critical applications, ensuring adherence to guidelines such as ethical decision-making in life-or-death scenarios. This could therefore represent a step forward in building ethically responsible mechanisms compliant with existing and future regulations and norms for autonomous systems.

The exact choice of decision-making structure and architecture depends on the context of the application and further investigating the various options would help understand the various possibilities and opportunities present. An initial exploration of such architectures was proposed in section 6.

7.5 Communication Optimisation

Intercession has the downside of increasing the amount of data needing to be exchanged to function. This is a direct result of the additional data involved along with specific design choices made to ensure the method performs as expected.

Note. *One instance of such a design choice is passing around a matrix containing all bids from agents for each task. Although this increases data transfer costs, it allows agents to always accurately determine the winning bid regardless of new ones received. As an example, an alternative approach could be to only store the current winning bid, its corresponding priority level, and target agent. This would however result in more messages being exchanged when an intercession occurs to replace a current winning bid with another weaker one (with a stronger priority). The system would still eventually converge on the correct allocation as each agent stores their own bids (c_{ijr} matrix), but the information would need to be re-propagated across the entire fleet as it is only stored locally. Instead, passing around all the bids placed for a specific task ensures that information only needs to be shared once to be taken into account regardless of the other's bids and bidding order. This way, the winning bid is always correctly determined upon having received all the bids once.*

To address these challenges, investigating how to efficiently manage the volume of information exchanged, and the rate at which to exchange it would ensure that the coordination protocol communication requirements remain managed and limited.

7.6 Advanced Prioritisation Mechanism(s)

The development of highly complex or permissive prioritisation mechanisms could further unlock the programmability and flexibility of the decision-making process. Developing prioritisation mechanisms going beyond the simple numerical scale proposed in this research to for example also incorporate digital signatures and/or more complex logic could significantly expand the application and properties of such mechanisms. Additionally, one may also consider dynamic prioritisation mechanisms, employing reputation-based approaches. As long as properties 1 and 2 are respected, any arbitrary prioritisation mechanism may theoretically be used.

7.7 Failure modes

It should finally be noted that while this coordination mechanism proves robust to communication failure, the same cannot be said about agent failure. In the event of an agent failing after having been allocated a task, two possibilities are considered:

- *The agent is unable to complete the task, but is still able to participate in the auction process:* In this case, the agent is still able to "report" its failure to the rest of the fleet through the emission of bids reflecting it. In other words, the agent emits a bid with a value of 0 for all the tasks it is incapable of taking any more (and/or simply stops bidding for new tasks).
- *The agent is **both** unable to complete the task and participate in the auction process:* This is the problematic case, as, in such a scenario, the rest of the fleet has no means of determining through the auction protocol that the agent is now unable to take on the tasks that were assigned to it. The lack of new bids emitted by the agent itself however at least ensures that no new tasks are assigned to it (note that the same does not hold for intercessions).

  Note. *This problem could be addressed through the introduction of additional logic in the prioritisation process, such as a heartbeat protocol leveraging the timestamps of the messages exchanged or the bids for failure detection. This however would need to be balanced to account for the possibility of the fleet simply experiencing a temporary communication failure.*

With all of the above points and comments having been established, it now becomes possible to consider the various applications of this principle to more concrete cases. More specifically, putting all of the above in the context of the SAR case study will help highlight the resulting dynamics and applications enabled by this approach.

8 Conclusion

This research's goal was to propose and investigate a novel mechanism applied to consensus-based auction algorithms, namely the process of *bid intercession*. To help effectively ground this work, we selected the *Search And Rescue problem* as the main MRTA case study to be analysed and used in the rest of this paper. The research then focused on establishing a clear understanding of the impact of the intercession principle on the underlying algorithms used. It is established that intercession effectively only modifies the scoring function used by the agents in the auction process. This in turn allows us to conclude that introducing *intercession in a consensus-based auction process does not impact key convergence properties of the underlying algorithm*, namely *convergence termination* and *convergence iteration complexity*. It is also demonstrated that intercession lets us *manipulate the outcome of a consensus-based auction effectively through controlling the source(s) of coordination and the bidding logic*. It was more specifically demonstrated that intercession allowed for running a fully centralised allocation process using decentralised processes as the main underlying mechanism. This along with additional demonstrations allowed us to conclude that the approach makes it possible to *manipulate a consensus-based auction outcome as needed while consistently maintaining the overall system's robustness*. All of the above claims were furthermore *verified experimentally*.

The research then focused on performing an *initial investigation of the novel decision-making architectures and dynamics enabled by intercession*. The One-to-Many, Many-to-Many, Many-to-One, One-to-One, and Combined architectures are proposed, along with respective use cases and applications in the reference Search And Rescue scenario. A few more concepts are explored and introduced, namely Restricted Auctions and a number of additional emergent dynamics resulting from the reorganisation of bidding responsibilities.

It can therefore be concluded that the intercession mechanism is not only a viable and appealing solution to certain MRTA problems, but it also paves the way to a large array of potential applications and extensions.

A large array of potential future research stems from these principles. First, all the components described in section 7 each present significant potential for improvements and enhancement to improve performances by themselves. Additionally, the coordination architectures described in section 6 are likely to hold additional prospects when considered within different contexts, and with different goals. Designing the correct coordination structure for the correct application therefore remains an area to be investigated.

References

[1] Luis Emmi, Mariano Gonzalez-de Soto, Gonzalo Pajares, and Pablo Gonzalez-de Santos. New Trends in Robotics for Agriculture: Integration and Assessment of a Real Fleet of Robots. *The Scientific World Journal*, 2014:e404059, March 2014. Publisher: Hindawi.

[2] Antonio Neves. *Service Robots*. BoD – Books on Demand, January 2018. Google-Books-ID: I8WPDwAAQBAJ.

[3] The Future of Robotics: Orchestrating the Heterogeneous Robot Fleet, May 2023.

[4] Boston Dynamics Partners With OTTO Motors To Coordinate Mobile Robots In The Warehouse.

[5] Brian P. Gerkey and Maja J. Matarić. A Formal Analysis and Taxonomy of Task Allocation in Multi-Robot Systems. *The International Journal of Robotics Research*, 23(9):939–954, September 2004.

[6] Felix Quinton, Christophe Grand, and Charles Lesire. Market Approaches to the Multi-Robot Task Allocation Problem: a Systematic Mapping and Survey. page 48.

[7] Jian Tang, Kejun Zhu, Haixiang Guo, Can Liao, and Shuwen Zhang. Simulation Optimization of Search and Rescue in Disaster Relief Based on Distributed Auction Mechanism. *Algorithms*, 10(4):125, December 2017. Number: 4 Publisher: Multidisciplinary Digital Publishing Institute.

[8] Navid Hooshangi, Ali Asghar Alesheikh, Mahdi Panahi, and Saro Lee. Urban search and rescue (USAR) simulation system: spatial strategies for agent task allocation under uncertain conditions. *Natural Hazards and Earth System Sciences*, 21(11):3449–3463, November 2021. Publisher: Copernicus GmbH.

[9] Guruprasad Airy, Tracy Mullen, and John Yen. Market based adaptive resource allocation for distributed rescue teams. April 2009.

[10] Jorge Peña Queralta, Jussi Taipalmaa, Bilge Can Pullinen, Victor Kathan Sarker, Tuan Nguyen Gia, Hannu Tenhunen, Moncef Gabbouj, Jenni Raitoharju, and Tomi Westerlund. Collaborative Multi-Robot Systems for Search and Rescue: Coordination and Perception, August 2020. arXiv:2008.12610 [cs].

[11] RoboCup Federation official website.

[12] Documentation – RoboCup Rescue Simulation.

[13] RoboCupRescue Robot League.

[14] Han-Lim Choi, Luc Brunet, and Jonathan P. How. Consensus-Based Decentralized Auctions for Robust Task Allocation. *IEEE Transactions on Robotics*, 25(4):912–926, August 2009. Conference Name: IEEE Transactions on Robotics.

[15] G. Ayorkor Korsah, Anthony Stentz, and M. Bernardine Dias. A comprehensive taxonomy for multi-robot task allocation. *The International Journal of Robotics Research*, 32(12):1495–1512, October 2013.

[16] Ahmed Hussein, Mohamed Abdelhady, Mohamed Bakr, Omar Shehata, and Alaa Khamis. Multi-robot Task Allocation for Search and Rescue Missions. *Journal of Physics: Conference Series*, 570:1–10, December 2014.

[17] Salvatore Aronica, Francesco Benvegna, Massimo Cossentino, Salvatore Gaglio, Alessio Langiu, Carmelo Lodato, Salvatore Lopes, Umberto Maniscalco, and Pierluca Sangiorgi. An Agent-based System for Maritime Search and Rescue Operations. volume 621, September 2010.

[18] Sven Koenig. The Power of Sequential Single-Item Auctions for Agent Coordination. page 5.

[19] Sahar Trigui, Anis Koubaa, Omar Cheikhrouhou, Habib Youssef, Hachemi Bennaceur, Mohamed-Foued Sriti, and Yasir Javed. A Distributed Market-based Algorithm for the Multi-robot Assignment Problem. *Procedia Computer Science*, 32:1108–1114, January 2014.

[20] P.B. Sujit and Randy Beard. Distributed Sequential Auctions for Multiple UAV Task Allocation. In *2007 American Control Conference*, pages 3955–3960, July 2007. ISSN: 2378-5861.

[21] Antidio Viguria, Ivan Maza, and Anibal Ollero. SET: An algorithm for distributed multirobot task allocation with dynamic negotiation based on task subsets. In *Proceedings 2007 IEEE International Conference on Robotics and Automation*, pages 3339–3344, April 2007. ISSN: 1050-4729.

[22] Luc Brunet, Han-Lim Choi, and Jonathan How. Consensus-Based Auction Approaches for Decentralized Task Assignment. In *AIAA Guidance, Navigation and Control Conference and Exhibit*, Honolulu, Hawaii, August 2008. American Institute of Aeronautics and Astronautics.

Appendix

8.1 Modeling Approach Overview

The modelling approach chosen breaks down the model into three key components: the environment, the agents, and the communication network. Each will serve a specific purpose in the simulation process:

- **Environment:** The environment provides information on the states of the various agents and tasks at a given task, along with details on the constraints and costs associated with the actions the agents can take. For this set of experiments, it was decided to use a grid environment, composed of nodes and edges. The resulting graphs are always connected but can be parameterised to include various ratios of missing to present edges. Additionally, each edge is randomly assigned a weight in a given range at generation.
- **Agents:** The agents represent the states and dynamics of the different members of a fleet. They each reflect the capabilities of different agent classes and possess the associated logic to interact with the environment and other agents. Two main agent classes will be used here; the base and the intercession agents. The two agents' categories are similar in all aspects but the bidding logic.
  - **Base agents:** The base agents will only bid for themselves and behave similarly to the agents in the original consensus-based approaches used as reference (CBAA/CBBA). They bid using the inverse of the **unweighted** Manhattan distance between their position and the target task position.
  - **Interceding agents:** The interceding agents possess the ability to bid for themselves **and** others. They bid according to the inverse of the **weighted** Manhattan distance.

  It is also important to note that the agents are kept updated on each other's states through state updates shared alongside the bid tables exchanged.
- **Communication network:** Although not necessary for this set of experiments, the communication network layer was implemented to emulate (if necessary) the constraints and dynamics imposed by the environment and the agents on the fleet's ability to exchange information. In the set of experiments considered here, however, it was assumed that the fleet possessed perfect communication. In other words, communication is not influenced by the environment and is reliable at all times. This is necessary to investigate the various hypotheses examined properly. More details on this is provided for each experiment in their respective description.

8.1.1 Simulator logic

The simulator used for this research was written from scratch to meet the requirements and constraints present in this work. The simulation leverages a ROS-like structure for the exchange of information and was entirely written in Python.

A number of goals were set to ensure that the simulator was able to correctly perform the various experiments necessary:

- Simulation of continuous time while not necessarily being able to run the simulations in real-time
- Simulation of continuous processes dynamics
- Control and simulation of differences in computational power

Two key aspects must be considered when looking at the functioning of the simulator employed. The management of the simulation and logical time, and the synchronisation of the different agents and processes with the simulation time.

- **Simulation and logical time:** To effectively control the simulation's progress, two distinct "times" are unidentified. Simulation time, which represents the amount of time elapsed during, and logical time, which breaks down the simulation time steps (referred to as an epoch here) into logical time steps. A simulation epoch consists of two logical steps (or phases):
  - **Consensus step:** The consensus step involves performing the exchange of messages and the coordination/task allocation process defined in each test. Here the agents do not evolve in the environment but only pass information around, and update their various local states based on the data received. The number of messages sent out by each respective agent during the consensus phase depends on the simulation time elapsed since the last message sent out, and the agent's message frequency.
  - **Environment step:** The environment step involves updating the states of the agents in the environment. During this step, agents use their local states to determine the move to perform to complete the various tasks assigned to them. The amount of steps undertaken by each agent depends on the agent's speed and

the amount of simulation time elapsed since the last move taken by said agent. It should be noted here that the agent's environment states do not constrain the other agents' moves. This means that the agents can take their respective environment step in any arbitrary order. For the sake of simplicity here each agent was given a speed of 1, meaning each agent only took a single step during each simulation epoch.

Figure 19 helps visualise the simulator time progression and management:

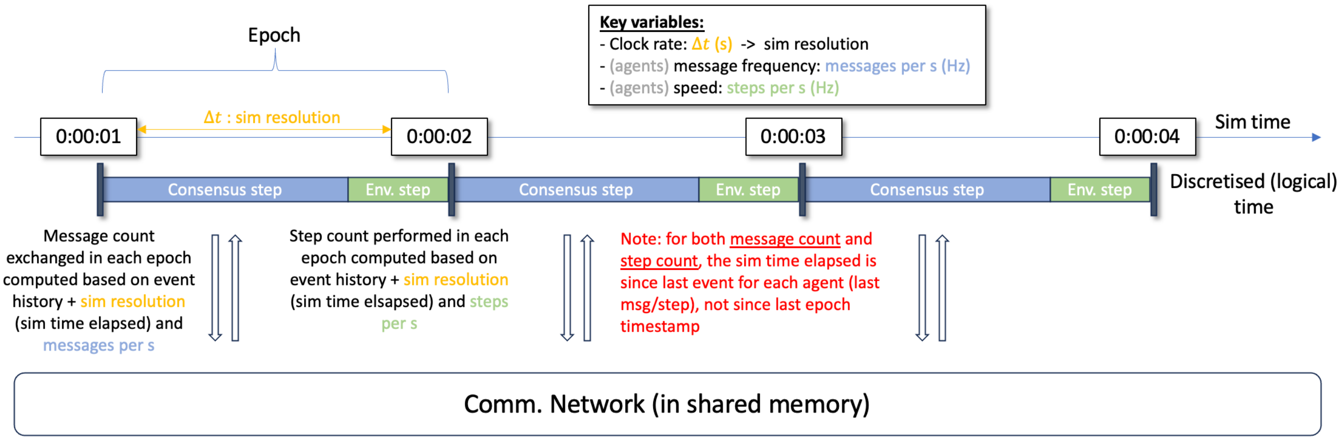

Figure 19: The simulation progress is broken down into simulation epochs, and logical sub-steps

- **Agents and processes synchronisation :** In order to capture the dynamics stemming from the agent fleet being a parallel and distributed system, it was chosen to run each agent in a separate process. This ensures that computations can be performed in parallel, ensuring a more true-to-life simulation process. This however still constraints the simulations to the computational speeds of the various processors used, making it hard to impose a predefined bidding rate for example. In order to compensate for the differences in computational time requirements of the various bidding mechanisms present across agents, the logical steps are used as checkpoints. Each agent process is made to wait for all other agent processes to have completed their computations before moving on to the next logical step. This ensures that not process/agent skips ahead of others, and is constrained in progress (in the simulation time) by the slowest process.

  A graphical representation of this is provided in Figure 20.

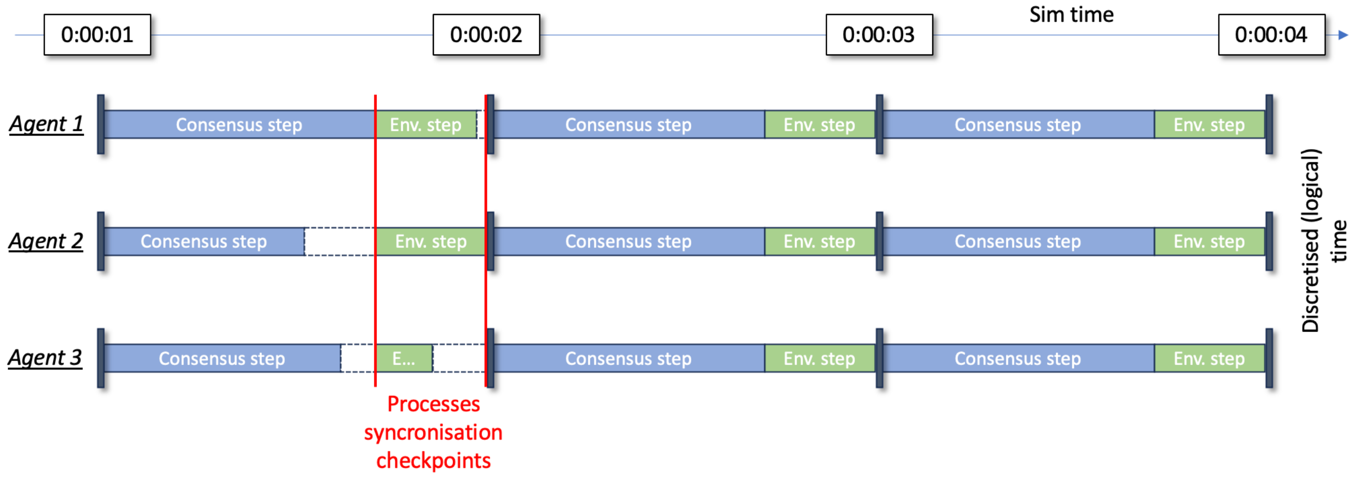

Figure 20: Three processes are shown here as an example. It can be noted that for each logical step, each process must wait for all others to have completed their respective logical step before moving on to the next one

This choice of breaking down the simulation process into the steps described above presents a large number of advantages. First, it allows for the precise analysis of the evolution and spread of information within an agent network (more on that in subsection 8.3). This additionally allows for simulating a continuous process by retaining the key dynamics which would impact the simulation results. In this case, it is the order in which the messages are exchanged, which directly affects the way information spreads within the network. Since all agents are synchronised in logical time (more in the next paragraph), we can also ensure that each agent correctly performs both a consensus and environment step at each simulation epoch, regardless of the computational cost associated with each respective step. This in turn effectively "erases" computational time and ensures that the specific processes (bidding notably) take place at the correct rate in simulation time irrespective of how heavy they are in comparison to other agents.

8.2 Experimental configurations

The table below (Table 1) recaps the main simulator parameters used for the different experiments. The details of the experimental setup themselves and the results are provided in separate sections below. It is important to note that a number of variables were kept constant across all experiments (unless otherwise stated), and as such do not appear on the table. Those are:

- Agents speed: 1 per s (one edge traversed)
- Message frequency: 1 per s

Table 1: Recap of the various important simulator parameters for the various experiments

	Experiment	**Exp. 1 (5.1, hyp. 3)**			**Exp. 2 (5.1, hyp. 4**			**Exp. 3 (5.2)**		
Variables		**Base**	**Min.**	**Max**	**Base**	**Min.**	**Max**	**Base**	**Min.**	**Max**
	Duration (s)	-	10	30	15	10	20	35	-	-
Sim	Task #	1	-	-	10	5	20	10	8	12
	Dual run	TRUE	-	-	FALSE	-	-	TRUE	-	-
	Const. agent #	TRUE	-	-	-	-	-	TRUE	-	-
Environment	Nodes #	-	-	-	60	45	75	60	-	-
	Connectivity %	-	-	-	0.8	0.6	1	1	-	-
Interceding agent(s)	#	1	-	-	1	-	-	1	-	-
Base agent(s)	#	5	3	8	5	3	8	5	3	8

8.3 Results: Exp. 1 - Allocation convergence experiment (subsection 5.1, hypothesis 3)

The allocation convergence experiment was designed in order to verify the conclusions attained in subsection 5.1 for the first hypothesis. Its goal is to verify the convergence properties of the proposed algorithm, and compare them to the ones of the original unaltered consensus algorithm (CBAA). To gain insight into the consensus states and convergence rate, the following approach was devised:

- Each test scenario will be run twice: once with, and once without intercession to compare the convergence properties of both every time.
- The speed of all agents will be set to 0. This is done to ensure that the bids do not change as the agents evolve in the environment.
- The message rate is set to 1 per time step in order to create a correspondence between time steps and logical progress of the algorithm.
- As explained earlier, the base agents bid using the Inverse Unweighted Manhattan formula as seen in Equation 5
- The most extreme role is selected for the interceding agent. It is tasked with interceding on behalf of all other agents (One-to-Many, see Figure 12). Additionally, the interceding agent also uses the Inverse Weighted Manhattan bid estimation (Equation 6), which is more computationally costly.
- For each run, a single task will be declared in the network. Agents will then exchange a single message at every time step and follow the task allocation protocol as expected.
- Upon completion of the run, the number of timesteps necessary to reach a consensus will be determined. Consensus is considered as reached once a single agent has been established as winner across all agents locally.
- The difference in time steps necessary is then computed to determine whether the introduction of intercession led to variations in the consensus process

The above experiment was run following a Monte-Carlo approach with the settings as shown in Table 1 in order to measure the convergence across a range of variables and configurations. The results were as follows:

Over a total of 93 runs, the average convergence rate for the consensus process without intercession was determined to be 1.75 time steps, whereas the average convergence rate for the same consensus rate with intercession was determined to be 2.13, leading to a difference of 0.38.

The results are as expected, with the small divergence being explained as follows. The number of steps necessary to reach consensus for this simulation are in the range of 1 to 2 (and occasionally 3) steps. This is a direct result of the stochastic nature of the message exchange process. To best explain this the process can be broken down as follows:

1. Initially, the agent which has discovered the task shares it with its neighbours.
2. Upon reception of the first message containing a reference to the new tasks, agents in turn compute their bid, which is then sent out in their next message

The process of sending out a message and receiving one are however not related, with agents sending out messages at a fixed rate, but processing the ones received upon reception. They are therefore performed asynchronously, and the ordering of the messages stochastic in nature (a desired dynamic in the simulator to best capture reality), which leads to variations in time step count to convergence.

It is possible to achieve convergence in a single time step if the message sent out by the agent (at time step t) is received and processed by all other agents before themselves sending out their own message for the same time step t. This is unlikely, resulting in our average being closest to the 2 time steps convergence. 2 time steps being required for convergence is, therefore, the most likely outcome (the agents only sending out their bids at time step $t + 1$. 3 time steps is the result of a flaw in the communication framework used. There is no way of ensuring that all messages sent out are received, and the publisher subscriber model used leads to messages being occasionally missed by an agent. A 3 step convergence is therefore an example of a case where an agent missed the initial announcement message, and was only made aware of the task at the second time step (leading it to possibly only sharing its bid at the 3rd step). This is however quite infrequent.

Finally, the slight bias in the average time step necessary for convergence between runs with and without intercession is a result of the extra computation time needed by the interceding agents. Unlike the base agents, the interceding agent computes bids for all other agents with a more complex path-finding logic, resulting in a larger delay before those may be computed upon reception of the message announcing the task. This makes the 1 step convergence practically impossible, as all agents (including the interceding agent) practically always send out their message before the interceding agent is done computing its bids.

It can therefore be noted that while the number of logical steps are unaffected by the introduction of intercession, the heavier bidding logic may lead to a slower convergence time. It is however solely a result of the heavier algorithms and not the mechanism itself. Assuming an ideal world where all bidding phases took the exact same amount of time to compute, the convergence rate would be the same for consensus-based auction processes with and without intercession (equal logical steps count).

8.4 Results: Exp. 2 - Allocation centralisation experiment (subsection 5.1, hypothesis 4)

This experiment was designed to verify whether it was possible to completely centralise a task allocation process which operated using consensus-based auction protocols. It is then proposed to perform a task allocation with a centralising interceding agent present and verify whether the final allocation corresponds to the one determined by said agent. The details of the simulator configuration can be found in Table 1. The centralising agent was given a priority level of 1, whereas all the base agents were given one with a value of 0.

The results were as expected, with the final allocation corresponding to the central interceding agent allocation every single time across the 201 runs performed.

8.5 Results: Exp. 3 - Allocation performance experiment (subsection 5.2)

The allocation performance experiment thrived to demonstrate the possibility of influencing the outcome of the auction process through the introduction of an interceding agent with a slightly more effective bid mechanism when considering the allocation evaluation metric. This experimental setup and its results are explained extensively in section 5, and as such will not be detailed here further to avoid repetition. The exact simulator configuration used can however be found in Table 1.

II

Literature Study

previously graded under AE4020

1

Introduction

The use of heterogeneous robots fleets has been growing in popularity in recent years across a variety of sectors, including warehouse management [1], logistics [6], and agriculture [15][38]. Although these systems are still in their early stages of implementation, they begin to gain momentum and exhibit promising potential across a wide variety of applications.

These fleets rely on effective communication and coordination to achieve robust and reliable task allocation, leveraging as efficiently as possible the knowledge and resources available to them. Additionally, robotic systems need to be built to function in cooperation with people [35] as human involvement is often important for activities that need on-site decision-making, complicated problem-solving, and adaptability. The challenge of coordinating and allocating tasks among multiple robots with varying abilities, also known as the Multi-Robot Task Allocation (MRTA) problem [18], is especially difficult and calls for effective communication and cooperation. In particular, the dynamic nature of the environments in which these fleets operate, the heterogeneous nature of the team member's individual skills, and the mission requirements pose considerable difficulty.

A notable application field for such systems is search and rescue [44] during a disaster. The search and rescue problem refers to the task of navigating, locating, and rescuing victims in a disaster-stricken area with an unreliable infrastructure (the problem can also encompass reacting to other emergencies such as fires and traffic obstructions). Multiple specialised actors must effectively coordinate to ensure the most efficient use of available resources is achieved to maximise the impact of the response. The tasks and goals are unknown initially, and evolve over time, making an initial planning phase impossible or simply ineffective. To further complexify the problem, the environment is only partially known, as the disaster may impact the infrastructure with various degrees of severity. This does not only potentially hinder movement, but also communication. Additionally, the unpredictable nature of most natural disasters makes them particularly hard to anticipate. This leads to the need for a system capable of being deployed rapidly, and operating in real-time and on-demand. The search and rescue problem is a difficult undertaking that calls for coordination and collaboration amongst numerous agents, including first responders, rescue teams, and specialised equipment. These agents may be mobile, sensing, or communicative, among other skills, and must constantly adapt to a changing and unpredictable environment. The search and rescue problem also includes the need to collect and process information about the affected area, effectively communicating information across the fleet given a limited and unreliable communication network, as well as identifying and prioritising potential victims for rescue. For those stranded in a disaster area, this is a crucial task that can mean the difference between life and death. To prevent further loss of life and property, competent search and rescue efforts are therefore essential.

Among the wide array of challenges presented by the search-and-rescue problem, task allocation proves to be particularly crucial in enabling agents and resources to best be leveraged to enable an effective response. The nature of the scenario however introduces a number of major constraints, which in turn must be accounted for when designing the various systems in charge of tackling this operation.

Many various approaches have been put forth to address these difficulties. Among them, auction-based strategies have had the most success (examples include [50][20][7]). Auction-based methods rely on virtual bidding and auctioning mechanisms to enable a group of actors to reach an agreement on the plan to be followed by the whole. These techniques particularly fall into one of two groups: distributed or centralised. Distributed approaches are more resilient but typically produce less optimal task allocation, whereas centralised methods typically produce better outcomes but are less resilient to communication disruptions. A novel method is proposed in this document, aiming to combine the best of both worlds to achieve a hybrid centralised-distributed method for task allocation in a dynamic environment for a heterogeneous team through the introduction of bid intercession in the auction process.

The objective of this literature review is to become familiar with the difficulties that arise in a search-and-rescue scenario as well as the cutting-edge tactics and methods that can be used to overcome those difficulties. Chapter 2 focuses on presenting the Search and Rescue problem, the associated problems and difficulties related to it, and the efforts that have been made so far to take on the challenge. A taxonomy of the Multi-Robot Assignment problem is then exposed in chapter 3, along with an analysis of the Search and Rescue problem using the provided prism. An introduction to the various existing market-based methods for tackling this problem class and it's variants is then provided in chapter 4, along with an introduction to the concept of bid intercession, which will be the main focus of this research. An overview of the concepts of Hierarchy And Reputation in multi-agent systems, a consequence of the new method proposed in chapter 4 is then given in chapter 5. A research proposal is then proposed in chapter 6, with the corresponding research methodology in chapter 7, and research planning in chapter 8.

2

Search And Rescue

The Search and Rescue (SAR) Problem is a highly complex yet important problem faced by society today. Search and Rescue situations can arise from a variety of causes, notably accidents, natural disasters, and other natural or human-driven events. In such situations, a large number of agents must coordinate in an attempt to locate any situations needing attention, overcome obstacles, and rescue lives and infrastructure as efficiently as possible. This task is further complexified by the fact that failing infrastructure may impact communication and movement capabilities, further hindering the ability of the search and rescue crews to effectively respond to emergencies.

2.1. RoboCupRescue

In order to better understand and frame this problem, the RobotCupRescue challenge can be used as a reference. The RoboCupRescue is part of the RoboCup organisation (Figure 2.1) created in 1996 [?]. RobotCup is an annual international robotics competition with as primary goal to promote robotics and AI research by offering appealing yet complex challenges for teams to take on [4].

Figure 2.1: The RoboCup was created in 1996 to promote robotics and AI

The competition involves more than 40 countries, and more than 3000 participants [2]. Their main goal of raising awareness of the various challenges associated with the various aspects of robotics is tackled by offering five different leagues:

- **RobotCupSoccer:** This league offers robotic soccer competitions, in which autonomous agents compete as humans would. The primary goal of this league is to further research in the areas of AI and multi-agent systems.
- **RoboCupRescue:** This league offers a framework for developing and testing systems to be used in disaster response scenarios.
- **RoboCup@Home:** This competition focuses on robots designed to assist with tasks in a home environment. Those can include cooking, cleaning, and more generally assisting humans when and where possible.

- **RoboCupIndustrial:** This league promotes research for industry-oriented applications. These include packaging, assembly, sorting, and any other applicable scenarios.
- **RoboCupJunior:** This league aims at promoting the field of Robotics with younger students, and promote learning and interest through providing a number of challenges to compete in.

2.2. SAR as Defined in RoboCupRescue

Focusing on the RoboCupRescue, the challenge offered by the league was designed to best reflect the reality of an emergency response and rescue during a disaster scenario. The league describes a natural disaster as a major adverse event. Those events impact the environment and infrastructure (notably transport and communication), resulting in a loss of shelter, food and supply shortages, and the potential spread of infectious diseases. Effectively responding to these events can ensure that economic losses and fatalities are minimised, and in turn speed up the recovery of the disaster-struck area.

Accordingly, four main objectives are therefore defined [2] to guide the emergency response effort:

- Save lives
- Prevent new disasters
- Collect data
- Short-/Long term planning

Similarly, the following limitation and constraints can be observed:

- Limited resources
- Incomplete information
- Real-time decision making
- Large number of actors involved
- Heterogeneity (in actors and tasks/task requirements)
- Dynamic scenario state

As such, a few key challenges resulting from the above goal and constraints are derived:

- **Victims detection:** The first challenge is the detection of human presence and vital signs in the disaster area, necessary to know what tasks need to be handled by the system. These then need to be located and relayed back to the emergency response team. This can be handled solely by the emergency response system, or/and can also leverage civilians on site to assist with relaying information back (though this also introduces new complexity).
- **Communication:** Developing a resilient communication system is absolutely critical, as ensuring stable communication is key to allowing for effective coordination and as a result an effective emergency response. This resilience can be achieved through better equipment, along with more resistant communication protocols and systems, capable of handling unreliable communication.
- **Task allocation:** Upon having established what tasks needed to be handled by the emergency response system along with their respective requirements, the agents must then coordinate effectively to ensure that these tasks are delegated efficiently and as optimally as possible to best leverage the resources present in the network (computational/manpower/equipment). This will be one of the main focuses of this paper's research.
- **Autonomous navigation/operations:** Upon having been assigned tasks, each agent must possess the ability to autonomously navigate to the desired/necessary location, and perform the necessary tasks related to its task (clearing up debris, extinguishing fires, delivering medical first aid). While some tasks are significantly easier for humans to perform, others require mechanical assistance. Limited manpower means that the more autonomous each robot will allow for a larger response capacity within a response team.

- **Others:** A number of other challenges can also be mentioned, such as the need for having quick-to-assemble robots allowing for scaling the supply to the demand as quickly as possible, or creating durable and resilient robots capable of surviving harsh environments resulting from the disaster.

2.3. Modelling SAR

The RoboCupRescue Simulator (RCRS) [2] shall be used as a reference frame to construct a model of the SAR scenario used in this research. Although the simulator itself might not be used during the experimental phase of this research (due to logistic reasons), it remains a solid reference of what a simulation for modelling such situations requires. It will as such remain a main driver for the structuring of the final framework used.

A number of key aspects are to be simulated to effectively represent the situation at hand. Those include the environment, the agents, and the communication network. More details on the representation of these different components is provided in the following sections.

2.3.1. Environment

Modelling the environment effectively is critical to ensure a representation of an emergency scenario is realistic enough to capture the complexity of the situation. The environment used in the 2023 edition of the competition can be seen below in Figure 2.2. It is an accurate representation of a part of Kobe, Japan.

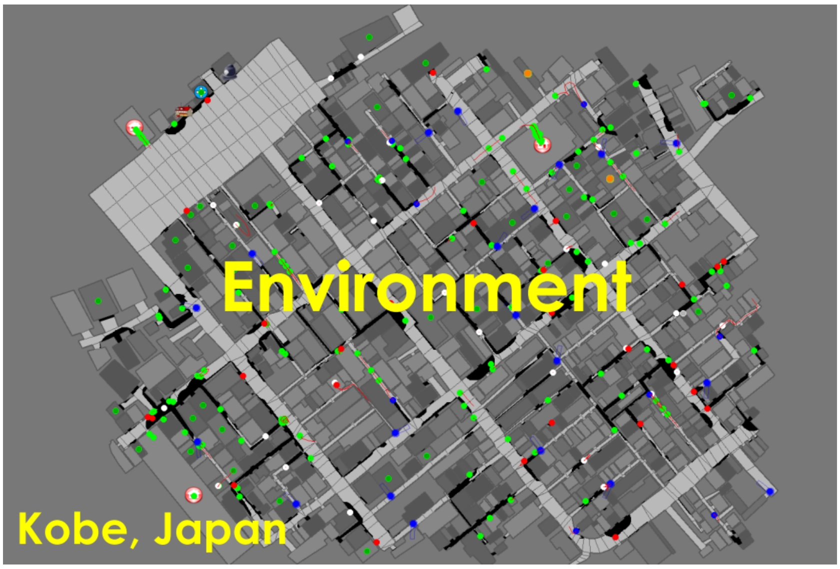

Figure 2.2: The environment used in the 2023 edition of the RoboCupRescue challenge

Figure 2.3: The RoboCup Rescue Simulator environment is capable of simulating the spreading of fire across buildings, and the collapse of infrastructure

The environment primarily acts as a baseline for agents described in subsection 2.3.2 to act on. It provides a position and information leveraged in turn by the communication simulation module (more on that in subsection 2.3.3).

It notably features buildings, roads, and various Points Of Interest (POIs) such as fire hydrants, gas stations, shelters, ambulance centres, police stations, and fire stations. The environment is capable of simulating the start and spread of fires of various intensities across buildings, along with the effect of firefighters turning off fires. Additionally, the environment can simulate road blockages, requiring the intervention of police forces to clear debris.

2.3.2. Agents

A total of eight classes of agents can be distinguished, each with its own set of properties and responsibilities. They can be seen below in Figure 2.4:

Figure 2.4: A total of 8 different agent classes are established in the RoboCup Rescue Simulator

- **Civilians**: The civilians are the most critical agents of the simulation, and play the role of victims in the disaster. They have varying attributes, such as age, health, size, etc. The civilians may require assistance from the different rescue teams and have as main goal to reach a refuge location.
- **Ambulance Team**: The ambulance teams are responsible for providing assistance and transporting injured civilians to refuges and shelters. The teams usually would comprise medical staff capable of administrating first aid, and limited medical supply and equipment. Ambulance teams rely on the ambulance centre to be dispatched and resupplied. Furthermore, ambulance teams rely on gas stations to refill the vehicles when necessary.
- **Fire Brigade**: The fire brigades are responsible for extinguishing fires across the disaster area, and rescuing civilians within burning and collapsed buildings. A fire brigade relies on water hydrant POIs to access water and the fire station for dispatching and equipment. Similarly to the ambulance teams, the fire brigades rely on gas stations to refill the vehicles when necessary.
- **Police Force**: The police forces' main role is to guide civilians toward shelters, and clear road debris. They rely on the police office for dispatching and equipment. They also rely on gas stations to refill their vehicles when necessary.

While all the agents described so far are mobile and on the field, the rest do not fulfil the same purpose. While the role of the above mobile agents is to intervene and operate on the field, the remaining static agents are primarily responsible for other tasks such as maintaining communication and coordination:

- **Refuge/Shelter**: The refuge is the safe haven for civilians and the end goal for each of them.
- **Ambulance Center**: The ambulance centre's main role is to coordinate and resupply ambulance teams. Information relevant to the various teams is gathered here and shared with the various teams to ensure an effective organisation. The ambulance centre is also responsible for coordinating with the fire station and the police office to orchestrate the emergency response.
- **Fire Station**: The fire station's main role is to coordinate and resupply fire brigades. To guarantee a successful organization, information pertinent to the various teams is gathered here and distributed to the teams. In order to coordinate the emergency response, the fire station centre must also work with the ambulance centre and the police department.

- **Police Office**: The primary responsibility of the police office is to coordinate and restock police units. To guarantee a successful organization, information pertinent to the various teams is gathered here and distributed to the teams. In order to coordinate the emergency response, the police department must also work with the fire station and the ambulance centre.

The above-mentioned agent must effectively communicate and coordinate to ensure that an effective response is provided when facing arbitrary scenarios. This is one of the reasons behind the RobotCupRescue organisers only releasing the actual scenarios used in the competition close to the competition date. This is done so to encourage the participants to develop systems either applicable to a wide range of scenarios and/or capable of being adjusted quickly to fit the case at hand.

2.3.3. Communication Network

The last critical component is the simulation of the communication network. Two main modes of communication are considered in the RCRS. Those are **voice** and **radio communication**. Agents are able to communicate directly with other agents if within a small distance radius. Otherwise, radio is available to the agents, at the cost of reliability (some messages could be lost/dropped, or simply delayed) and certain restrictions (the communication channels operate on a limited bandwidth).

While the exact communication simulation process is not specifically explained in the simulator's documentation, a number of possible implementations [14] and factors [47] are possible.

Those notably include ray tracing, which implies simulating the path radio waves would follow (bouncing around and going through various materials) to account for the obstacles encountered. While a full ray tracing simulation reflecting the physics of radio waves would prove most accurate, it is also the most costly computationally. It is furthermore very complex to effectively simulate the reflectivity/absorption of all the material present, making it a partially attractive solution for large-scale simulations. Lighter alternatives involve tracing the direct path between the two agents communicating and assigning to all obstacles on the said path a probability of impacting communication, along with properties such as how the communication would be affected. While slightly less accurate, this method is significantly cheaper. A final possible implementation is simply to couple (inversely proportional) the probability of messages going through with the distance between the emitter and receiver. This is the least accurate method but by far the cheapest one when computational power is a constraint. Many other factors can also be accounted for (such as the strength of the signal emitter for example), and a viable trade-off between computational cost and accuracy can therefore be achieved effectively for this component.

2.4. Optimising Task Allocation in Search and Rescue

The following research will focus specifically on the task-allocation aspect of the SAR problem. This is chosen as the ONERA expressed a particular interest in this area. This specific problem is known more generally as the Multi-Robot Task Assignment problem (MRTA) and will be discussed in more detail in chapter 3. The research will aim at tackling the SAR problem with a particular focus on the communication challenges present. This most notably include the existing communication constraints and failures. Other aspects such as the difficulty associated with different tasks requiring different agents to be effectively addressed will therefore be ignored.

3

Multi-Robot Assignment Problem

Task allocation in multi-agent robotics systems is a major challenge in unlocking more sophisticated and autonomous transportation systems. The problem, known as the Multi-Robot Task Allocation [18] (MRTA) problem, consists of assigning a set of tasks to the various members of a robotic team. Understanding the exact nature of the MRTA problem at hand is critical to understanding the underlying complexities and limitations of the operating scenario. As such a reference taxonomy of the MRTA problem is initially discussed in section 3.1 bellow. It is used in turn in section 3.2 to frame the Search and Rescue MRTA problem and clarify the challenges associated with it.

3.1. MRTA Taxonomy

A number of attempts were made to establish a taxonomy of the possible problems at hand. Notable taxonomies include the one presented in [18], which proposed a breakdown of the problem centred around three key aspects with two modalities respectively. Those are (in no particular order):

1. **Robots' capabilities:** Robots may be able to take on only a single task (**ST**), or multiple tasks at once (**MT**)
2. **Tasks types:** Certain tasks can be achieved by a single robot (**SR**), while others allow for/require multiple robots (**MR**) to cooperate/coordinate together
3. **Assignment and Constraints:** Some tasks must be taken on solely based on the instantaneous knowledge available (**IA**). Others on the other hand allow for a time-extended knowledge (**TA**) to be provided. This item was extended in [39] to also include a further distinction between tasks with time-window constraints (**TA:TW**) and those with synchronisation constraints with other tasks (**TA:SP**)

This definition was then extended a first time in [32] which added a distinction for the different dependencies across tasks and robots:

4. **Degree of inter-dependence:** Problems in which the order of execution does not impact the value of tasks are denoted as No Dependencies (**ND**). Tasks for which the order of execution impacts the tasks' values are referred to as In-schedule Dependencies (**ID**). If the order of execution depends on other robots, the problem is qualified as Cross-schedule Dependent (**XD**). Finally, the problem is categorised as Complex Dependencies (**CD**) if the specific decomposition selected to take on a task affects the eventual value of the task.

The taxonomy was enriched once more in [45] to include a categorisation of the communication constraints present along with the level of uncertainty present in the environment and tasks:

5. **Communications:** The communication categorisation is made up of two elements:
    - **Connectivity Type:** The communication is considered as Global (**G**) if all robots are able to communicate directly with all other agents present. If this is not the case, the

communication is then said to be Local (**L**).

- **Uncertainty upon messages:** The communication can be either certain (**N** for no uncertainty), with no risk of message losses, or lossy (qualified as **S** for stochastic).

6. **Environment and Tasks:** This categorisation aims at specifying the dynamics (or lack thereof) present in the problem environment. It is also made up of two elements, one for the A priori Knowledge, and the other for the environment itself:
   - ***A priori* Knowledge:** Tasks may be either known (**K**) before execution, or unknown (**U**)
   - **Dynamics:** The environment and tasks present in said environment may be either static (**S**), or dynamic (**D**) in nature.

The full taxonomy and possible categories are summarised in item 3.2.

3.2. The SAR MRTA Problem

With the MRTA taxonomy established, it is now possible to analyse the challenges of task allocation applied to the Search and Rescue scenario. Looking at each element of the above-described taxonomy one at a time, the following problem qualification can be established:

1. **Robots' capabilities - MT:** The agents defined in our scenario are heterogeneous in nature. While the ambulance teams, fire brigades, and police forces would most likely qualify as a single task (**ST**), all the coordinating centres could have the ability to take on multiple communication tasks simultaneously, making the problem multi-task (**MT**)
2. **Tasks types - MR:** The tasks possible in this scenario range from simply picking up an injured civilian to coordinating multiple fire brigades to effectively control a fire (requiring or simply allowing multiple agents to effectively contain and extinguish). The scenario tasks are as such qualified as a combination of single (**SR**) and multi-robot (**MR**) tasks.
3. **Assignment and Constraints - TA:TW and SP** The taxonomy yet again highlights the sheer complexity of the scenario as all assignment and constraint categories can be observed in the various tasks. Tasks such as picking up a civilian who called (and then hung up) for help would fall under the instantaneous knowledge available (**IA**). This is different from the scenario where the person stays on the phone to keep the centre up to date on the progress of a fire for example (which would in this case fall under **TA**). Additionally, a fire must be handled within a specific time (the window starts when the fire is detected, and ends when the building reaches a critical state) resulting in a time window constraint (**TA:TW**). Finally, rescuing a civilian may only be possible once a specific road blockage has been cleared, resulting in a task synchronisation constraint (**TA:SP**).
4. **Degree of inter-dependence - CD:** The scenario could yet again span a large number of categories given the large variety and nature of tasks. To start off, the choice of orders for which civilian to pick up or which fire to deal with first has an impact on the result of the other civilian/firefighting tasks, leading to In-schedule Dependencies (**ID**). Simple tasks such as clearing blockages are not time-sensitive, but they may be necessary for other tasks taken on by different agent classes to be made achievable, resulting in Cross-schedule Dependencies (**XD**). Finally, the chosen allocation and order of task completion will have a direct impact on the overall solution quality (implying complex dependencies **CD**), making the problem **CD** overall.
5. **Communications:**
   - **Connectivity Type - L:** For the case of a natural disaster, it is assumed that the communication network is unreliable, making the communication local (**L**)
   - **Uncertainty upon messages - S:** For the same reasons as above, the communication is assumed as lossy/stochastic (**S**)
6. **Environment and Tasks:**

- ***A priori* Knowledge - U:** None of the tasks would be known before the actual event, making this scenario a **U** (it can be noted that some areas may be known for being more vulnerable, allowing for an "expectation" of where tasks might appear).
- **Dynamics - D:** The environment and tasks are dynamic (**D**) as a new task could appear at any point (a fire starting somewhere), or disappear (residents managing to turn off a fire), and the collapse of roads and buildings could change the environment at any point.

As can be seen in the above breakdown, the Search and Rescue scenario generally presents enough complexity to fall into most categories simultaneously when it comes to attempting to classify the tasks and tasks properties. Additionally, the communication is simply assumed to be unreliable, severely limiting the potential methods employable. Finally, the dynamic nature of the environment and scenario results in the need for a very resilient and dynamic protocol capable of re-organising itself efficiently in real-time.

Degree of Inter-Dependence	**Task Type**	**Robot Type**	**Assignment and Constraints**	**Communications**		**Environment and Tasks**	
No: ND	2=Single-Robot: SR	2=Single-Task: ST	Instantaneous: IA	1=Connectivity Type	1=Uncertainty upon messages	1=A priori Knowledge	Dynamics
In-schedule: ID			Time-Extended: TA	1=2=Global: G	2=No uncertainty: N	1=2=Known: K	2=Static: S
Cross-schedule: CD	2=Multi-Robot: MR	2=Multi-Task: MT	w/Synchronisation and Precedence: TA:SP	1=		1=	
Complex: XD			w/Time Windows: TA:TW	1=Local: L	Stochastic: S	1=Unknown: U	Dynamic: D

4

Market-Based Allocation Methods

The actual task assignment mechanisms can differ significantly between implementations, and numerous strategies have previously been theorised and tested in an attempt to tackle the MRTA problem class. Among the different techniques developed, the Market-Based Approaches (**MBA**s) have proven particularly successful and versatile (attempts targeted to the SAR specifically are described in section 4.3).

Market-Based Approaches primarily consist of providing robots with a simulated economic market in which to trade tasks and resources in an attempt to solve the Multi-Robot Task Allocation Problem [45]. Given a reward function and a cost function, robots calculate their own utility associated with each task available, which is then utilised to carry out the actual task assignment process. MBAs have proven to be very popular due to their efficiency, flexibility, and ease of implementation. Many variations have been considered and tested, mostly striking a balance between a few crucial characteristics depending on the application problem under consideration. They are the process's degree of centralisation, the task allocation rate, and the bid mechanisms. When attempting to categorise the approaches, two broad groups may be found:

- **Auction-based methods:** Those methods are highly centralised and rely on an auctioneer to perform the final allocation

- **Consensus-based approaches:** Those methods rely on a consensus mechanism to perform the task allocation, avoiding a single point of failure at the expense of a potentially more optimal solution.

They are further discussed in section 4.1 and section 4.2.

4.1. Auction Based Methods

The first and most often used MBA strategies are auction-based ones. In those plans, robots act selfishly and generally coordinate using a greedy strategy. Such procedures are often centralised, with a single auctioneer in charge of the process administration, supervision, and decision-making.

The Single-Item (SI) auction [31] is the most basic illustration of an auction-based MBA. In this method, the auctioneer serves as the principal coordinator, and the process is centred around it. The fundamental idea is shown bellow in Figure 4.1:

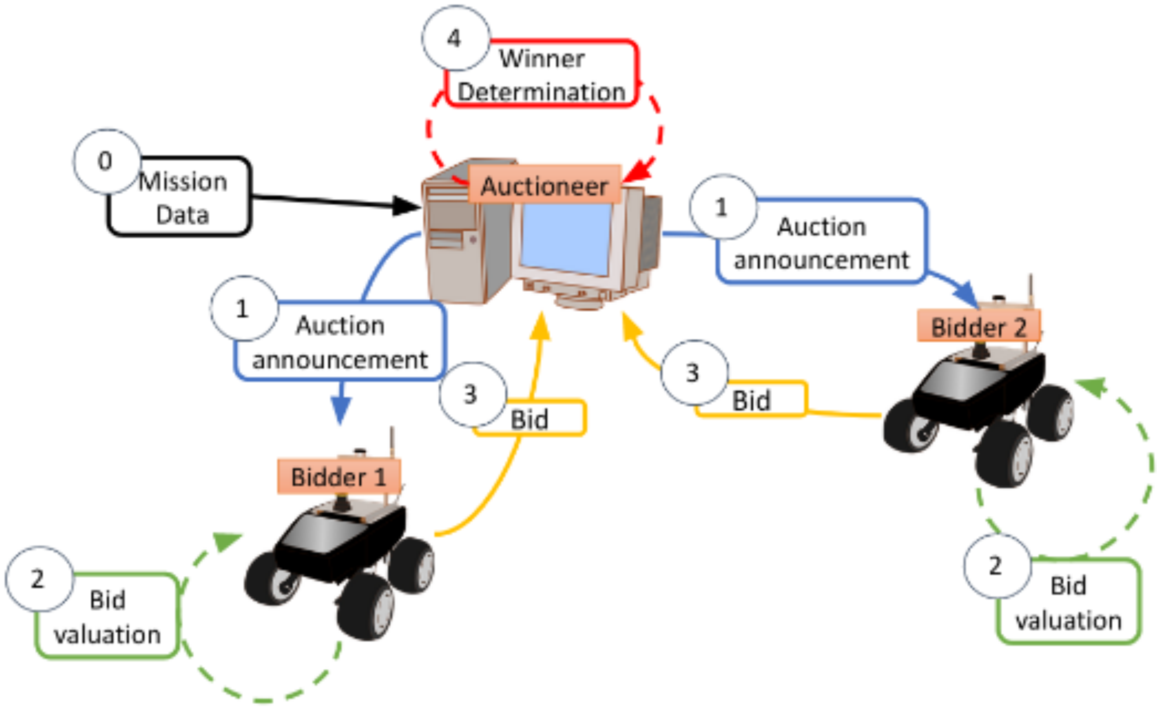

Figure 4.1: The SI auction process [31], an example of Auction-based methods (figure from [45])

0. The auctioneer is given a list of tasks to be auctioned off at the start of a mission.
1. A task is then chosen by the auctioneer and declared to be "for sale," often via an announcement message.
2. Each robot taking part in the auction computes a "bid" for the task that is being auctioned off after receiving an announcement message. This "bid" is dependent on a number of properties. Any combination of information, including the robot's location in relation to the work, the amount of battery life left, the robot's aptitude for the task, etc., may be included in this.
3. The robots then relay their bid to the auctioneer.
4. The auctioneer evaluates all of the proposals and assigns the task to a robot in accordance with them after receiving all of the bids or after waiting for a predetermined period of time.
5. If any tasks remain unsold, the procedure is then repeated.

A number of key sub-categories can be noted from the Auction-based MBAs. The main ones along with their respective properties are explained in the subsections below.

4.1.1. Notable Variants

The **Single Item Auctions** [31] (SI) is the most elementary market-based allocation method, and is the one described in the steps above. To further improve the allocation speed, **Parallel Single Item Auctions** [51] (PSI) auction and distribute all tasks in a single round. This is however possible if tasks are completely unrelated and do not influence each other. If that is not the case and the order of tasks taken on is important, **Sequential Single Item auctions** [49] (SSI) operates in such a way that agents take the set of tasks in their current plan into account when computing their bid. Finally, a **Combinatorial Auctions** [52] gets agents to bid on all combinations of tasks possible to then achieve the most effective bundle allocation. A few more alternatives revolving around this allocation mechanism exist, but most simply introduce modifications to the core logic to better address a specific scenario challenge or aspect listed in the MRTA taxonomy.

All the above-cited methods come with trade-offs, both in terms of optimally, communication loads, and computational power requirements. It as such becomes necessary to match and pair allocation algorithms with their corresponding MRTA problem presenting the right properties.

It should be noted that all the above methods rely on a central coordinating instance, making them all non-resilient to communication failure. As such they most likely would not prove effective in the SAR given the scenario requirements.

4.2. Consensus-Based Approaches

The consensus-based approaches still leverage the same bid and auction mechanisms but distributed the decision-making process to allow for a more resilient system. By coupling an auction phase with a consensus phase, consensus-based approaches adopt a more decentralised strategy [11, 13]. By doing so, the singular point of failure flaw found in auction-based methods can be avoided during the task allocation process. A process for the Consensus-Based Bundle Algorithm, a consensus-based task allocation algorithm, is depicted below in Figure 4.2:

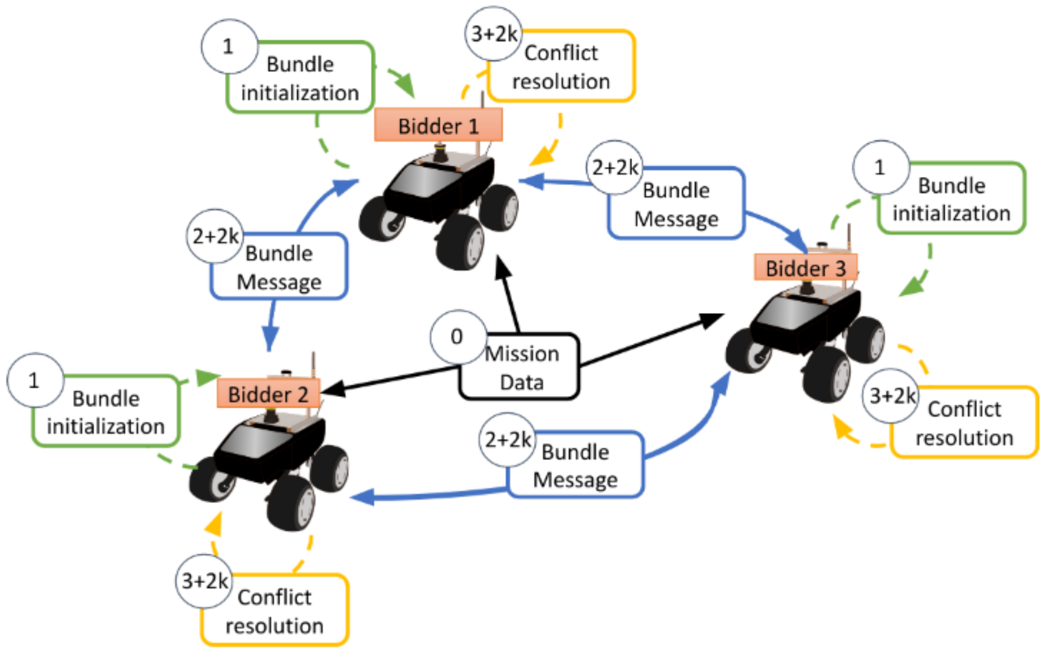

Figure 4.2: A simple consensus-based auction process: CBBA [11], more details below (figure from [45])

1. Each robot initialises a bundle of tasks during the auction phase, adding tasks to its bundle in decreasing order of preference.
2. Each robot then transmits to the other robots its bundle with its corresponding preferences values (representing the robot's bids for each job).
3. After getting the bundle messages, robots look for conflicts, which are tasks in their bundles that are also in a teammate bundle.
4. If there are any, they move on to the conflict resolution phase, where the bids are compared to decide which robot should forego the contentious tasks. A robot will abandon a given task and all tasks following it in its bundle if another robot successfully outbids it on the said task. The robots then reconstruct their bundles without taking the tasks they were outbid on into account.
5. Robots then share their respective bundles once more (step$(2 + 2k)$ with k=1), which could lead to their peers entering a new phase of conflict resolution (step $(3 + 2k)$ with $k = 1$). The consensus loop continues (and k is incremented) until a conflict-free allocation is attained.

The Consensus-Based Auction Algorithm and Consensus-Based Bundle Algorithm serve as the foundation for most adaptations of this strategy (both discussed in the subsections below). The fundamental appeal of these classes of methods is the trade-off of some optimally (CBAA and CBBA guaranteeing 50% provided some specific assumptions; see [11][13]) for a substantially more resilient decentralised process.

4.2.1. Consensus Based Auction Algorithm (CBAA)

The Consensus-Based Auction Algorithm (CBAA)[13] is a simpler variant of the case described above. The major difference being the fact that agents exchange individual tasks instead of bundles, and simply compare the bid for tasks individually instead of within the context of a bundle.

It is worth noting that despite the tasks being allocated individually, the allocations themselves can (and do) occur in parallel. While being comparably less performant than CBBA with respect to allocation quality, CBAA therefore has the major advantage of requiring fewer message exchanges to reach a consensus. This in turn allows for the algorithm to generally converge faster.

Going back to the MRTA taxonomy discussed in chapter 3, the main constraints of this method find their roots in the task-by-task nature of the algorithm. It as such prevents the approach from covering most of the **Assignment and Constraints** spectrum. Similarly, all of the possible categories of **Degree of inter-dependence** are outside of the scope of this algorithm (with the exception of the **ND** category).

4.2.2. Consensus-Based Bundle Algorithm (CBBA)

The Consensus-Based Bundle Algorithm (CBBA)[13] is the one depicted in Figure 4.2 and the corresponding description. It operates similarly to CBAA but is expanded to the multi-assignment problem, where each agent is given a sequence of tasks to complete.

This results in a consensus-based method capable of better incorporating constraints and relationships present within a set of tasks in the allocation process, yet again covering a larger part of the MRTA taxonomy spectrum.

It is important to note that CBBA is substantially different that the Combinatorial Auctions mentioned earlier in subsection 4.1.1. While all task sequence combinations are evaluated in Combinatorial Auctions, the auction process is performed at the task level in CBBA. This trade-off is chosen to balance convergence speed and communication/computation loads. It will generally allow for a more effective allocation to be found when compared to CBAA, but achieves this through a comparably higher overall consensus cost. It however tends to remain the most common baseline for the creation of variants which extend CBBA to cover a larger scope of the MRTA taxonomy mentioned in.

4.2.3. Other Variants

Most consensus-based approaches variants use the CBBA algorithm as foundation, with each variant focusing on a different property subset of the MRTA taxonomy. Notable ones include the **Asynchronous Consensus-Based Bundle Algorithm** [25, 26] (ACBBA), which extends the base algorithm to allow for asynchronous operations, critical given modern communication protocols and their respective drawbacks. **Coupled-Constraint Consensus-Based Bundle Algorithm** [55] (CCBBA) extends CBBA to enable accounting for coupled constraints to address more complex characteristics. **Team Consensus-Based Bundle Algorithm** [8] (TCBBA) nests CBBA coordination processes to enable allocating tasks to multiple teams efficiently. The concept of relays and relay tasks are introduced in **CBBA with Relays** [42] to help maintain a unified communication network topology. **CBBA with IRRT** [43] aims at unifying a motion planning algorithm (Information-rich Rapidly-exploring Random Trees) with a CBBA planning phase for an efficient data collection and fleet coordination process. Additionally, [21] considers tasks requiring several agents, and [41] focuses on multi-mode tasks (tasks which can be solved in multiple different ways) with also require multiple agents. Finally, the Performance Impact Algorithm proposed in [54] is a modified version of CBBA, which uses a different bid evaluation method referred to as "significance", and improves on a few of the consensus rules for conflict resolution.

Many other variants have been proposed, each focusing on enhancing different aspects of the base CBAA and CBBA algorithms. As of today, most of the research and interest seem to be focused on consensus-based variants over auction-based approaches, with most recent papers published focusing on CBBA enhancements and variants [45].

Finally, a number of other papers propose radically different approaches, which still leverage consensus mechanisms without necessarily relying on auction methods. Notable examples include [24], which proposes a distributed Hungarian algorithm, or [40], which leverages a distributed genetic algorithm to find effective solutions. Both methods cited here also produced a few variants such as [33], [36].

4.3. MBA applied to the SAR MRTA problem

MBA has proven to be a popular choice for tackling the MRTA aspect of the SAR problem. The principles are leveraged by [50] to perform a local task allocation using centralised bidding mechanisms, where upon discovering a task the agent respectively takes on the role of auctioneer, and auctions the task off. A multi-step and responsibility method is described in [20], which proposes to delegate the role of prioritising tasks, and selecting auctioneers to a central agent, therefore producing a multi-level coordination structure with multiple central elements. A centralised approach for decision-making is also considered in [7], [22], and [9] (the latter focusing on search and rescue scenarios in maritime environments).

A few papers also investigate the use of distributed approaches, such as [56], which focuses on dynamic task allocation, and [40].

While all the methods presented above have their merits, they clearly highlight a split of the various approaches into two categories; those that focus on performance (without worrying about communication failure), and those that focus on reliability. A common pattern can be associated to this observation. Methods aiming at maximising performances use centralised approaches, while those focused on resilience adopted more decentralised systems. The two are therefore analysed in the next section.

4.4. Centralised vs. Decentralised Communication Architectures

The reliance on a centralised versus a decentralised coordination architecture represents a critical distinction between the two market-based allocation categories mentioned above. The two costliest operations in an auction-based method are the bid estimation, estimated for each agent bidding in a round, and the winner determination process (WDP). While for both categories described above the agent bid estimation is performed similarly by each respective agent, the WDP methodology is done differently. For centralised systems, data is collected in a singular "location" by the auctioneer to compute a solution, which is then dispatched back to the network. This is in stark contrast to the decentralised architecture, which relies on local knowledge and rules for performing the allocation and achieving consensus. This divergence in ethos has significant implications on the requirements and resulting use cases.

For centralised architectures, key notable advantages include the simplicity/convenience resulting from having all data gathered in one place. This allows for well-established optimisation algorithms to be leveraged (such as MILP and meta-heuristics), given the auctioneer performs the task allocation based on a global view of the system. Examples of that include Cao et al's approach, which leverages particle swarm optimisation (PSO) algorithms to solve the MRTA problem [12]. This additionally makes it possible to concentrate processing power around the actor requiring it, allowing for optimising the auctioneer agent with better computing specs, at the expanse of other capabilities, such as displacement and sensing. This in turn allows for more computationally expensive algorithms such as meta-heuristics or MILP to be used, allowing for more optimal solutions to be found. Accordingly, the global view provided by the centralisation of data makes it theoretically possible to find globally optimal solutions.

The main negative aspect of centralised architecture is the presence of a single point of failure. All agents must be able to communicate with this central agent, and a failure to do so leads to an effective system failure. This dependency on reliable communication introduces a significant number of constraints, such as on range, reliability, and any other factors impacting communication performance. Additionally, a centralised system scales as long as the central agent is capable of coping with the extra computational load, which can be limiting.

Decentralised solutions specifically aim at addressing this issue by removing the central element characteristic of centralised approaches. Consensus on task allocation is achieved solely through peer-to-peer exchanges, ensuring resilience to an unreliable and heterogeneous communication network. Every agent effectively takes on the role of bidder and auctioneer, and the overall auction is split into numerous sub-auctions, allowing for a decentralised consensus to be achieved efficiently and reliably. The decentralised nature of the algorithms also allows for much more effective and reliable scaling, as while the eventual solutions found can be less optimal, the systems are still capable of reaching local consensuses, ensuring that the network keeps moving and does not fail.

These types of approaches however come at the significant cost of giving up the prospect of attaining optimality in the allocation, as a global view of the system states is never considered.

4.5. Hybrid Approach, a Research gap?

While significant research has gone into investigating the above communication architecture, very little can be found on hybrid architectures. A number of approaches split the allocation process into separate steps, which are then done sequentially, partially centrally and partially decentralised [12]. No papers could be found however proposing methods that seamlessly combine centralised and decentralised mechanisms, making it a clear research gap to investigate.

Each approach was designed for a specific MRTA subset of problems, and performs a tradeoff between resilience and performance by design. This was done so to ensure the problems of applications' various properties could be tackled, ensuring a robust process when necessary. Robust however does not necessarily imply optimality, and while a decentralised process for example allows for effectively withstanding a dynamic communication network topology, a well-built network might not encounter those communication problems frequently. It might the majority of the time operate as a fully connected network, with communication failure only occurring a small fraction of the time. In such a scenario, the edge case of a failing network forces the adoption of a decentralised process, resulting in a sub-optimal process the majority of the time. No solutions to this tradeoff exist currently in the current literature.

Upon careful analysis of the different existing approaches, both centralised and decentralised, a proposal was devised which aims at blending the two ethos, and lets both architectures operate simultaneously together, allowing for leveraging the advantages of both when relevant/applicable. This proposal, referred to as **bid intercession**, is detailed in section 4.6 below.

4.6. Centralised Operations in a Decentralised Setting: Bid Intercession

While centralised and distributed processes tackle communication differently, both categories of approaches make the same assumption when it comes to the source of the bids. All methods and papers reviewed for this literature study operated under the assumption that each agent was always responsible to compute and emit a bid for themselves. Changing this assumption can be done through the introduction of the **bid intercession mechanism**, the main contribution and focus of this research.

Bid intercession is defined as the capability of agents to compute and place bids for tasks on behalf of other agents. This process has the advantage of integrating almost seamlessly within the CBAA and CBBA mechanism, only requiring some logic to select which bid to keep when performing the auction process. While less impactful for centralised processes, the introduction of this capability in distributed processes opens up the systems to many new dynamics and network behaviours, which will be the main subject of this research. Three features in particular will be focused on:

- **Centralised processed emulation:** The main advantage of the introduction of bid intercession in a distributed approach is the resulting (theoretical) ability of a distributed network to effectively operate as a centralised one when the communication network topology allows for it. One such example would be through an agent bidding for all other agents, and its bids always being considered as the main bid in the auction process (an initial bid estimation could still be performed locally by all agents, which would in turn be considered by the "central" agent performing the WDP before emitting a fresh set of bids reflecting the desired final task allocation). In such a system, given a fully connected network, the above-mentioned agent would effectively act as the centralised agent, resulting in a centralised decision-making process. This would in turn allow the system to naturally transition from a more optimal centralised architecture to a distributed one as the communication network evolves.
- **Positive asymmetry augmentation:** Another significant advantage of the introduction of bid intercession is the possibility for a much more heterogeneous fleet to be assembled, presenting significantly larger gaps in capabilities within a single network. This is mainly due to the fact that bid intercession partially lifts the on-the-edge constraints and requirements, as agents do not necessarily need to be deployed on the field and take on tasks to be able to participate in the task

allocation process. This means that it now becomes possible to include in a fleet a data centre for example, with as sole purpose to compute optimal bids based on the information relayed to it. This in turn greatly increases the positive asymmetries possible in a given network (a positive asymmetry (PA) is a statistical term describing a value distribution which is not symmetrical around the mean).

- **External/human control in the decision-making process:** The introduction of bid intercession opens up interesting avenues for allowing external control within an autonomous distributed system. Human intervention through the placement of prioritised bids would allow for interfering and "piloting" the decision-making process without impacting the fundamental mechanisms in place.

Another important aspect of this concept is knowing when to intercede in a network. The answer to this question may vary heavily based on the situation and desired behaviour. In a scenario where bid intercession is used to enable humans to interfere in the task allocation process, the intercession would be only occasional and used when intervention is desired. In contrast, a more systematic intercession could be established (with a single central agent continuously bidding on behalf of all the other agents) when centralised behaviours are desired. This would allow for the network to retain the ability to remain functional and autonomous in the event of a loss of communication, while still making use of a more optimal central planner when possible. Finally, a more dynamic approach could be considered, where agents intercede on behalf of others when deemed necessary. This could be a result of agents having a poor reputation when it comes to estimating bids, or simply because certain agents know that their estimations are more accurate (better sensors, more up-to-date data, etc). The exact intercession dynamics can as such be parameterised to obtain the desired allocation behaviours, further allowing for "parameterising" the task allocation process.

4.6.1. Key Technical Challenge: Information Source Discrimination

Such systems however raise a number of important challenges, and most notably results in a need to discriminate between different information sources when conflicts arise as a consequence of intercession.

When an agent α intercedes by placing a bid with value V on behalf of another agent β for a task w, the fleet finds itself having to choose which bid to consider for agent β when performing the auction for task w:

$$\begin{aligned} B_{\beta:\beta\to w} &= V_1 \to \text{bid emitted by agent } \beta \text{ for itself} \\ B_{\alpha:\beta\to w} &= V_2 \to \text{bid emitted by agent } \alpha \text{ on behalf of agent } \beta \text{ (intercession)} \end{aligned} \tag{4.1}$$

It is important to note that in the above case: $B_{\beta:\beta\to w} \neq B_{\alpha:\beta\to w}$

When such conflict arises, an extra data point becomes necessary to ensure that the whole fleet considers the same bid (and consequently performs the same auction) despite the process operating in a distributed fashion. This need to discriminate information sources can therefore be addressed by introducing some elements of hierarchy and/or reputation. Finally, if all advanced discrimination methods fail, a fall-back rule must be included, ensuring a tie-break is always provided. Such fallback rule or metric could be as simple as comparing the IDs of the conflicting bid's owners (and retaining the one with the highest index), or the time of the bid emission (with the earliest bid getting priority).

These concepts are reviewed in chapter 5 before describing the proposed research and methodology for investigating in detail the impact of intercession on a distributed auction process.

5

Hierarchy And Reputation In MAS

Hierarchy and reputation are the two main discriminatory elements that can be considered when evaluating the importance and reliability of an information. Those concepts prove particularly important within the context of bid intercession, where agents find themselves having to select a source of information to follow when multiple are available. The concept of hierarchy and its various possible declinations are first discussed in section 5.1. The idea of reputation is then addressed in section 5.2.

5.1. Hierarchy

Hierarchy is a first possible approach to discriminating between different information sources in the event of a conflict. It can also be seen as a form of **hard trust**, as the rules within the group are fixed.

A hierarchy can be defined as a system or organisation in which individuals or things are ordered, ranked, or classified according to their importance, authority, or responsibility. A hierarchy contains different levels or tiers, and each level has its own set of properties, privileges, and responsibilities. The higher levels of the hierarchy are usually more important and dominant than the lower levels. Additionally, a hierarchy can also be implemented using a set of rules, allowing for more complex systems to be put in place, distributing roles and liabilities in a less linear fashion. It is assumed here that the hierarchy levels are known and observed globally, meaning all agents follow the same hierarchical structure.

Hierarchies are immutable, and their non-dynamic nature allows for easy scaling to large and complex systems, along with the possibility of nesting various hierarchical structures etc. Implementations may vary from simply ordering each agent in order of priority, to more advanced relative rankings accounting for the agents' various capabilities and performance. Such hierarchies must however be created with all participating actors in mind to avoid conflicting definitions and conflict resolution.

5.1.1. Applicability and Limitations

Providing a level of "priority" for each data point (setting a hierarchy) can prove to be an effective and simple solution to the challenge of information source selection. Given that the priority levels do not collide, the prioritisation is unambiguous, ensuring an error-free convergence respecting the hierarchy in place regardless of how distributed the network is. This approach proves applicable in the scenario where the priority in the source of information is clear and easy to establish. Those include:

- **Human intervention:** In the event of human intervention, it could be necessary to treat the instructions as absolute. This could easily be achieved by making the human-produced bid the highest priority possible.
- **Superior agents:** It could be that an agent with superior sensing and computational power is present in the network, with as main role to intercede and provide highly accurate bids without taking on any tasks. In such cases, and assuming the agent is always capable of producing a more accurate bid estimation, it could simply be provided with a higher hierarchy level, ensuring prioritisation of its contribution.

Adopting a hierarchy, or priority-based system does however come with significant drawbacks. It is to start with absolute in nature and does not account for the quality or confidence of the information. An agent with a high hierarchy level will always be prioritised, no matter how incorrect (or how frequently incorrect) its involvement is. This inability to evolve and self-adjust could prove detrimental to the overall performance of the network over time, and introduce inefficiencies. Additionally, this could prevent certain further developments, such as the introduction of mechanisms allowing the system to optimise itself based on its own performance.

As such, while hierarchy-based conflict resolution proves invaluable and necessary in certain specific cases, finding a more flexible and dynamic solution might be necessary to best leverage all information available and produced in the network.

5.2. Reputation

Reputation is a key notion in a multi-agent system. It describes the general impression or perception that one agent has about another agent or group of agents based on their previous encounters and observations. One way to conceptualise reputation is as a type of "social capital" agents utilise to make choices and choose which other agents to interact with [3]. Going back to the previous section, reputation could be considered as a form of **soft trust**, or dynamic hierarchy.

An agent's reputation in a multi-agent system is often expressed as a score or numerical value that reflects the agent's previous interactions and behaviour. Other agents frequently assess the agent's reliability or trustworthiness using this score. A high reputation score, for instance, increases the likelihood that other agents would trust and work with that agent, whereas a poor reputation score may make other agents wary or avoid interacting with that agent.

The kind of observation the agents will use to form a judgement about other agents can be divided into two main categories:

1. **Direct observations**: Agents watch other agents in action and utilise what they learn to alter how they perceive their reputations. The agents must have access to a lot of information about one another in order for this strategy to be effective.

2. **Indirect observations**: Agents base their judgement of an agent's reputation on the assessments and opinions of other agents. For example, an agent may ask other agents for feedback on a particular agent's behaviour or use online rating systems to evaluate the reputation of other agents.

It should be noted that the above two categories are not mutually exclusive, and an agent may estimate the reputation of another agent on the basis of a combination of direct and indirect observations.

A number of motivating questions can therefore be considered when investigating reputation in multi-agent systems:

- How can decision trust be represented?
- How can reputation be represented?
- How can evidence (observations) be connected to reputation/trust?
- How should uncertainty be handled (such as when an agent's reputation is rarely met)?
- How should the transitivity of reputation be handled sensibly?

5.2.1. Formalisation

To better grasp the extent of the challenge, the following formalisation [3] can be considered to better frame the problems and challenges at hand. A few definitions must first be established:

- **Agents:** The agents are the actors responsible for taking on the tasks, and rating each other. In this example, the agents are referred to as α, β, and γ.

- **Predicate** X**:** The predicate is the item or piece of data being considered in the evaluation

- **Trust** $T_{\beta \to} X$ (in a predicate)**:** The trust an agent β has in a given predicate X (Figure 5.1).

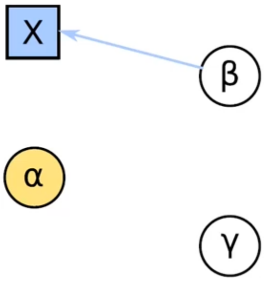

Figure 5.1: Trust, also referred to as **decision trust**, as it measures the security of a given agent acting on a given predicate [3].

- **Reputation** $R_\alpha \rightarrow \beta$ **[0, 1]:** The reputation is the trust an agent α has on another agent β (Figure 5.2).

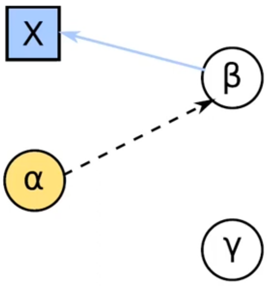

Figure 5.2: Reputation, also know as **reliability trust**, is a measure of the expectation a given agent (α) has that another agent (β) will interact reliably [3].

A simple example of the above would be the α and β being the various actors in the SAR network, and X is a predicate about the bid of an agent for a given task.

Finally, a distinction must be made between global and local reputation. In the event that $R_{\alpha\rightarrow\beta} = R_{\gamma\rightarrow\beta}, \forall\gamma \neq \beta$, the reputation is considered global. All agents share a common reputation space and belief, then referred to as R_β. If this does not hold, the reputation is then considered local. Within the context of distributed task allocation systems (considered for this work), the reputation is categorised as local, as the agents do not share a common global state.

5.2.2. Overview of State of the art

Four main reputation systems stand out from the literature. Those are:

- The Naive Binary Rating
- The Beta Reputation System
- The Dirichlet Reputation System
- Subjective logic
- Others (fuzzy logic, etc)

Naive Binary Rating (NBS) is the simplest of all systems mentioned above. It relies on prior binary evaluations of a set of interactions to determine the maximum likelihood of an interaction being positive. In such a system, the rating is modelled as a Bernoulli variable q_ϑ, representing the predisposition of an agent ϑ to behave in a desirable way. Given

$$\begin{aligned} n_g(t) = n_g &\text{ is the number of } \textbf{positive} \text{ ratings up to time } t \\ n_b(t) = n_b &\text{ is the number of } \textbf{negative} \text{ ratings up to time } t \end{aligned} \tag{5.1}$$

The likelihood (Equation 5.2) and maximum likelihood (Equation 5.3) are defined as

$$L\left(q_{\vartheta}\right)=\mathbb{P}\left(n_{g}, n_{b} \mid q_{\vartheta}\right)=\binom{n_{g}+n_{b}}{n_{g}} q_{\vartheta}^{n_{g}}\left(1-q_{\vartheta}\right)^{n_{b}} \tag{5.2}$$

$$\widehat{q_{\vartheta}}=\arg \max_{q}(L(q))=\arg \max_{q}\left(n_{g} \log(q)+n_{b} \log(1-q)\right)=\frac{n_{g}}{n_{g}+n_{b}} \tag{5.3}$$

The Naive Binary Rating approach only depends on the observations, and therefore corresponds to the naive intuition that "an agent that behaved correctly in the past will continue to behave well, and as such deserves a higher reputation". While seemingly simple in principle, this rating system has proven particularly important, notably given the fact that humans tend to binarize ratings [17] by mainly considering the extremes when asked to rate on a scale. A number of different methods and variants leveraging this approach have been proposed, including for example methods for distributed peer-to-peer networks [46]. Research was also put into investigating the robustness of such systems [30] in an attempt to gauge their objectiveness and reliability.

The **Beta Reputation System (BRS)** [23], while somewhat similar to NBR, is a model based on Bayesian Network and beta probability function which attempts to go a step further by also integrating an a priori (Figure 5.3) in the reputation estimation process. This a priori is defined per application, ensuring an initial tailoring of the reputation system fitting a given application/scenario.

Similarly to NBR, the rating is modelled as a Bernoulli variable q_{ϑ}, representing the binary set of interaction outcomes $\{\text{good}, \text{bad}\}$. The Bernoulli parameter is defined as a random variable itself. From there, two probability density functions (PDFs) are obtained, the **a priori**, and the **a posteriori**. Both belong to the Beta probability density functions family, with hyperparameters α_0 and β_0 encoding the prior knowledge of the a priori PDF (how likely an agent is for being right/wrong at the beginning), and $\alpha(t)$ and $\beta(t)$ the information of the a posteriori PDF.

The a priori PDF is defined as follows (Equation 5.4):

$$f_{q_{\vartheta} \mid \alpha_{0}, \beta_{0}}(q)=\frac{\Gamma\left(\alpha_{0}+\beta_{0}\right)}{\Gamma\left(\alpha_{0}\right) \Gamma\left(\beta_{0}\right)} q^{\left(\alpha_{0}-1\right)}(1-q)^{\left(\beta_{0}-1\right)} \tag{5.4}$$

$$0 \leq q \leq 1, \quad \alpha_{0}>0, \beta_{0}>0$$

The same convention as NBR (seen in Equation 5.1) is adopted here for the observations annotation, $n_g(t)$ for the good interaction count, and $n_b(t)$ for the bad interaction count. From there, given the posterior hyperparameters $\alpha(t)=\alpha_0+n_g(t)$ and $\beta(t)=\beta_0+n_b(t)$, the a posteriori PDF is defined using Equation 5.5:

$$f_{q_{\vartheta}}\left(q \mid n_{g}, n_{b}\right)=\frac{\Gamma(\alpha(t)+\beta(t))}{\Gamma(\alpha(t)) \Gamma(\beta(t))} q^{\alpha(t)}(1-q)^{\beta(t)}=f_{q_{\vartheta} \mid \alpha(t), \beta(t)}(q) \tag{5.5}$$

The reputation of an agent at an instant t is then represented using the a posteriori PDF of q_{ϑ}, which expresses the likelihood of agent ϑ behaving honestly. To obtain a scalar to use in various algorithms, the expectation of q_{ϑ} is often used (Equation 5.6)

$$R_{\alpha \rightarrow \vartheta}(t)=\mathbb{E}_{q_{\vartheta} \mid \alpha(t), \beta(t)}[q]=\frac{\alpha(t)}{\alpha(t)+\beta(t)}=\frac{\alpha_{0}+n_{g}(t)}{\alpha_{0}+\beta_{0}+n_{g}(t)+n_{b}(t)} \tag{5.6}$$

While still relatively simple, this approach allows for incorporating more information into the reputation mechanism, allowing for a pre-calibrated state to be included. The model was later extended by the original authors to support partial satisfaction, forgetting factor, and two extra operators for combining and discounting opinions [10] (see the Dirichlet Reputation System below). The BRS forms the basis of a large number of reputation-based systems nowadays, such as sensor-integrity management

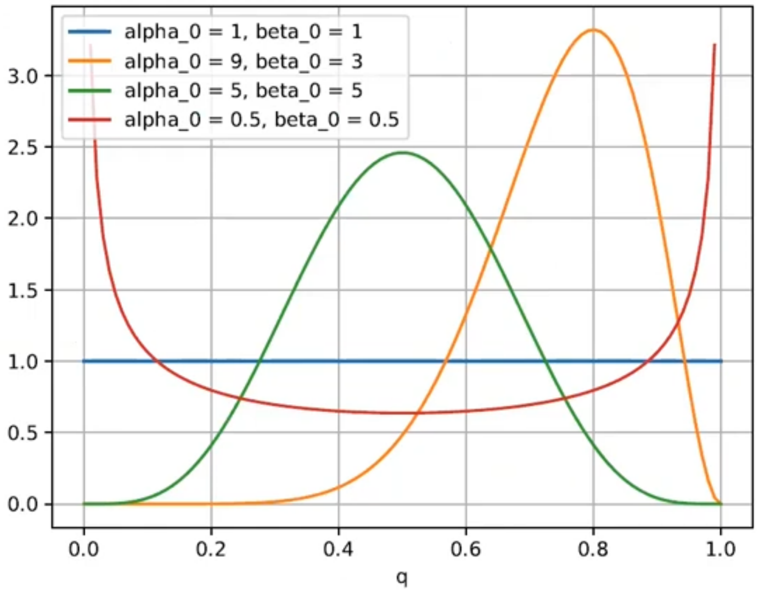

Figure 5.3: The a priori PDF [3] is used to encode in the distribution how likely an agent is of being right or wrong. It ensures a correct initial calibration of the rating system when known initially

[16], reputation-based protection schemes for distributed networks [53], and other possible applications [37].

The third reputation system considered is the **Dirichlet Reputation System (DRS)** [27]. DRS operates based on similar principles to the Beta Reputation System, but uses Dirichlet distributions (a multivariate generalisation of the beta distribution, see Figure 5.4). This in turn allows these systems to incorporate an element of how good (or how bad) an interaction was, enabling more nuance in the observations by modelling multivariate probability. Similarly to the BRS, an a priori PDF is used to "prime" the reputation state, and an a posteriori PDF is used to obtain the reputation of a given agent at a given instant in time.

For this model, the rating is characterised by K parameters q_{ϑ_i}, such that $\sum_{i=1}^{K} q_{\vartheta_i} = 1$. The vector q_ϑ is therefore as a K-variate random vector. From there, the a posteriori PDF is obtained using Equation 5.7:

$$f_{\boldsymbol{q}_\vartheta|\boldsymbol{\alpha}_0}(\boldsymbol{q}) = \frac{\sum_{i=1}^{K} \Gamma\left(\alpha_{0i}\right)}{\prod_{i=1}^{K} \Gamma\left(\alpha_{0i}\right)} q_i^{(\alpha_{0i}-1)} \tag{5.7}$$

$$0 \leq q_i \leq 1, \quad \sum_{i=1}^{K} q_i = 1, \quad \alpha_{0i} > 0$$

The observation vectors are then made of K elements all equal to 0, with the exception of the one corresponding to the rating, which is set to 1. The ratings can then simply be collected by adding up all the vectors. Ponderation with respect to time (ageing) may also be included through performing a weighted addition. This is done so using a longevity factor $\lambda \in [0, 1]$, which controls how discounted an observation is based on its age (the more outdated an observation is, the less impact it has on the final reputation rating).

From there, a number of possible representations can be considered. Those include:

1. **Evidence representation:** This is the most straightforward representation and simply involves expressing the aggregate evidence vector $\vec{R}_y$. The amount of ratings of level i for agent y is then denoted by $R_y(i)$.

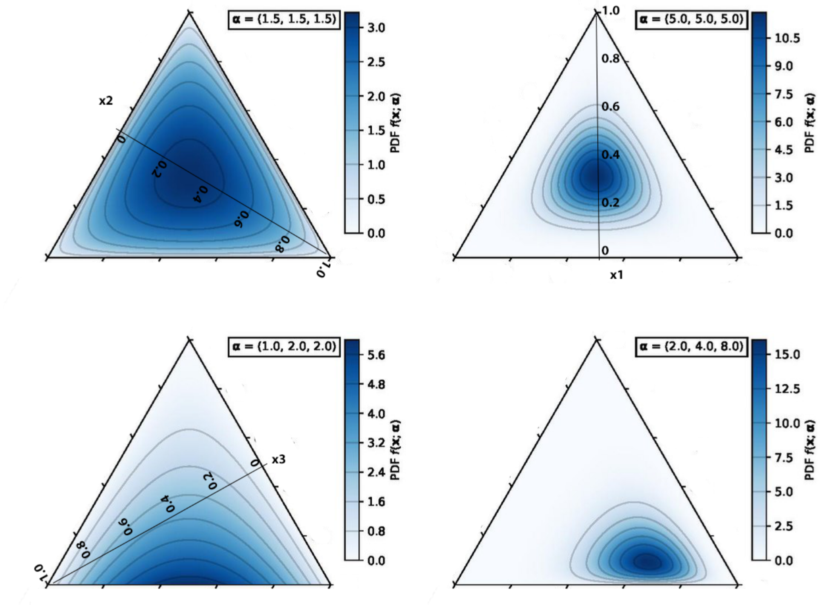

Figure 5.4: The Dirichlet distribution [5] is a generalisation of the beta distribution, and allows for modelling continuous multi-variate probability

2. **Density representation:** This is the PDF obtained using Equation 5.7.
3. **Multinomial representation:** This representation involves defining the reputation score as a function of the probability expectation value of each element in the state space for an agent y. The expectation value corresponding to each rating value can be determined using Equation 5.8, resulting in a vector $\vec{S}$:

$$\vec{S}_y : \left(S_y(i) = \frac{R_y(i) + Ca(i)}{C + \sum_{j=1}^{k} R_y(j)} ; \mid j = 1 \dots l \right) \tag{5.8}$$

The reputation score $\vec{S}$ can then be interpreted as a multinomial probability measure of how a particular agent would be expected to behave in a future interaction.

4. **Point estimate representation:** The reputation score can also be summarised as a single data point. This can be achieved by for example assigning a point value v to each individual rating level i, and solving for the normalised weighted point estimate score σ (Equation 5.9).

$$\begin{aligned} \nu(I) &= \frac{i-1}{k-1} \\ \sigma &= \sum_{i=1}^{k} \nu(i) S(i) \end{aligned} \tag{5.9}$$

While a single point of comparison can prove particularly convenient for certain applications, this representation however leads to significant losses in information, such as the polarisation of the ratings.

DRS is an effective rating system, which significantly improves on BRS while retaining most of what makes BRS effective. It should be noted that the algorithm is sensitive to the initial base rates (a priori) provided, and while enough observation could eventually compensate for an incorrectly calibrated baseline, the algorithm would take longer to correctly converge.

The last major approach is the **Subjective Logic** system [29] [28]. The key principle behind subjective logic is to separate and explicitly express probability estimates and uncertainty. The key unit of information in subjective logic is the Opinion, a mathematics object representing belief and uncertainty. An opinion is made up of four key components:

$$\text{Opinion: } \omega_x = (b, d, u, a) \tag{5.10}$$

1. **Belief - b:** The belief mass supporting x as being True (direct result of evidence)
2. **Disbelief - d:** The belief mass supporting x as being False (direct result of evidence)
3. **Uncertainty - u:** The belief mass uncommitted
4. **Base rate - a:** The a priori belief, used when no committed belief mass is present

$$\begin{gathered} b \in [0,1] \quad d \in [0,1] \quad u \in [0,1] \quad a \in [0,1] \\ b + d + u = 1 \end{gathered} \tag{5.11}$$

Opinions can be derived using the previous reputation methods described, such as BRS and DRS PDFs. The formalism for manipulating opinion objects is referred to as Subjective logic (see Table 5.1).

Table 5.1: A correspondence between probability, set, and logic operators [3]

Subjective logic operator	Symbol	Binary logic/ set operator	Symbol	Subjective logic notation
Addition	$+$	Union	$\cup$	$\omega_{x\cup y} = \omega_x + \omega_y$
Subtraction	$-$	Difference	$\setminus$	$\omega_{x\setminus y} = \omega_x - \omega_y$
Multiplication	$\cdot$	AND	$\wedge$	$\omega_{x\wedge y} = \omega_x \cdot \omega_y$
Division	$/$	UN-AND	$\overline{\wedge}$	$\omega_{x\overline{\wedge}y} = \omega_x / \omega_y$
Comultiplication	$\sqcup$	OR	$\vee$	$\omega_{x\vee y} = \omega_x \sqcup \omega_y$
Codivision	$\overline{\sqcup}$	UN-OR	$\overline{\vee}$	$\omega_{x\overline{\vee}y} = \omega_x \overline{\sqcup} \omega_y$
Complement	$\neg$	NOT	$\overline{x}$	$\omega_{\overline{x}} = \neg\omega_x$
Deduction	$\circledcirc$	MP	$\parallel$	$\omega_{Y\parallel X} = \omega_X \circledcirc \omega_{Y\mid X}$
Abduction	$\overline{\circledcirc}$	MT	$\overline{\parallel}$	$\omega_{Y\overline{\parallel}X} = \omega_X \overline{\circledcirc} \omega_{X\mid Y}$
Discounting	$\otimes$	Transitivity	$:$	$\omega_x^{A:B} = \omega_B^A \otimes \omega_x^B$
Cumulative Fusion	$\oplus$	n.a.	$\diamond$	$\omega_X^{A\diamond B} = \omega_X^A \oplus \omega_X^B$
Cumulative Unfusion	$\ominus$	n.a.	$\overline{\diamond}$	$\omega_X^{A\overline{\diamond}B} = \omega_X^A \ominus \omega_X^B$
Averaging Fusion	$\underline{\oplus}$	n.a.	$\underline{\diamond}$	$\omega_x^{A\underline{\diamond}B} = \omega_x^A \underline{\oplus} \omega_x^B$
Averaging Unfusion	$\underline{\ominus}$	n.a.	$\underline{\overline{\diamond}}$	$\omega_X^{A\underline{\overline{\diamond}}B} = \omega_X^A \underline{\ominus} \omega_X^B$
Belief Constraining	$\odot$	n.a.	$\&$	$\omega_X^{A\&B} = \omega_X^A \odot \omega_X^B$

Opinions once derived can therefore be manipulated and chained up to determine reputations within and across a multi-agent network, leveraging the beliefs of all agents and their confidence levels. This proves particularly useful within the context of distributed multi-agent networks but also introduces significant complexity. The flow of reputation within a multi-node network was studied in a number of papers [48], [34], yet it remains a complex field possibly worth investigating further, particularly for task allocation processes.

A large number of other reputation methods have been proposed, and are valid possible solutions for the problem at hand. A lot of them are recapped in [19], which proposes an overview of the various methods and paradigms up to 2015. The literature review also breaks down in more detail the various elements that may be considered when determining the reputation within a multi-agent system, and the various categories of cooperation possible within a single network.

5.2.3. Applicability and Limitations

The above-described reputation estimation approaches may prove effective to deal with the novel complexity introduced by bid intercession, but could in turn also have a significant impact on the underlying task allocation mechanisms. A stand-alone reputation system does not necessarily result in a global consensus on the reputation of every agent in a network, which proves to be problematic within the context of consensus-based task allocation.

Consensus-based systems are capable of reaching an equilibrium given that all agents follow the same rules. While potentially more optimal, the introduction of a reputation system effectively introduces a second consensus process, which turns into a dependency/constraint for the overarching task allocation process. It is impossible to ensure a consensus is reached on task allocation as long as a consensus on the source of truth (which bid to consider, a.k.a the reputation/hierarchy of the information) is not found within a given network.

Being able to set up a mechanism ensuring that consensus can effectively be reached on the reputation would however allow for a network to self-moderate and optimise at run time, potentially significantly improving its own performance in light of the data collected during operations.

6
Research Proposal

A preliminary literature study suggests that while a number of methods solving Multi-Robot Task Allocation Problems exist, each is very tailored to a given scenario, and performs a tradeoff between resilience and performance by design. Furthermore, all market-based methods considered until now operate under the premise that each agent is responsible for placing their own bids, irrespective of the task allocation process.

A literature gap is found on what could be referred to as "bid intercession", and its impact on distributed consensus-based task allocation methods. Bid intercession enables agents in a market-based allocation system to place bids on specific tasks for other agents, a concept never explored before from what was observed in existing literature. This would potentially open up new dynamics within a distributed coordination system, allowing for hybrids of centralised and distributed coordination algorithms to be developed. Bid intercession has the potential to allow for the creation of fail-safe decision-making systems, capable of leveraging all available resources while ensuring that the process remains robust to communication failures, allowing for a reliable and flexible system in uncertain dynamic environments.

As such, a research objective is first presented in section 6.1, with a corresponding research question and sub-questions in section 6.2.

6.1. Research Objective

Upon having investigated the topic extensively, and in light of the literature survey presented earlier, the following research objective is proposed:

> "To investigate the implications of introducing bid intercessions in a multi-hop, peer-to-peer communication network for distributed task allocations through simulations and real-world experimentation."

The research is therefore practice-oriented, and aims at investigating the impact of an alteration to the current state of the art, in an attempt to improve over existing methods. The driving goal of this research will be to make an attempt to propose a "one-solution-fits-all" method for tackling the MRTA challenge, capable of performing well over a wide range of network configuration and allocation objectives/constraints while retaining fail-safe properties. The research will focus specifically on the distributed class of task-allocation algorithms described in section 4.2 earlier, which will act as the baseline for the development and testing of bid intercession.

6.2. Research Questions

To effectively tackle the research objective proposed in section 6.1, the following research question is formulated:

> "How is a ***consensus-based distributed task allocation process*** in a ***dynamic uncertain***

scenario impacted by the ***introduction of bid intercession*** in the ***multi-hop, peer-to-peer, dynamic communication network*** of a ***distributed heterogeneous agent fleet***"

To best tackle the main research question, a number of key questions were established, each aimed at addressing a specific major component of the research. Three key areas of investigation were defined:

1. The impact of the introduction of intercession on the dynamics/mechanisms
2. The impact of the introduction of intercession on the allocation performance/quality/optimality
3. The level of external control on the allocation process enabled by the introduction of intercession

Three questions were constructed to address those, which are in turn further broken down into sub-questions, each addressing a particular hypothesis on each aspect of the topic. Finally, an approach is devised for investigating each hypothesis:

SQ-1: What are the functional properties of a consensus-based task allocation method with bid intercession

This question addresses the need to understand the impact of the introduction of bid intercession in the task allocation processes and mechanisms. Two main aspects are considered, notably the impact on the underlying supporting algorithms, and the resulting novel flexibility in distributed decision-making process architectures. Three driving sub-questions are established here:

SQ-1.1: Does bid intercession have any convergence impact on the underlying distributed methods?

Convergence is possibly the most elementary requirement for consensus-based protocols. It is fundamental to ensuring networks eventually reach an agreement over any decision process. As such it is the first aspect investigated in this research. Bid intercession does not modify any of the underlying mechanisms when introduced. It is simply an injection of information, with some additional logic for establishing which source of information to consider. The choice of coordination system adopted for information source selection (hierarchy/reputation) therefore becomes the focal point here. To specifically address the impact of bid intercession of intercession the following hypothesis is proposed:

H-1.1 The process of bid intercession does not have any impact on the convergence of the underlying distributed method **given that the source of truth is unambiguously defined at all times**. More specifically, the following are not impacted:

- Ability to converge
- Rate of convergence

The main rationale behind this hypothesis is that demonstrating that the introduction of bid intercession does not impact the underlying algorithm mechanisms allows for inheriting the underlying algorithm's convergence guarantees.

Validating this hypothesis would therefore answer the sub-question. This would be demonstrated first theoretically, and confirmed experimentally through the creation of a number of toy problems to use as a baseline to help isolate the desired cases best proving (or disproving) the hypothesis (more in chapter 7).

SQ-1.2: What are the keys aspects/challenges introduced by bid intercession

The goal of this question is to dive deeper into the fundamentals of the bid intercession logic and focus on detailing the challenges and concepts resulting from its introduction in a decision-making process. This question, which will be addressed theoretically, mainly focuses on the distributed selection of a source of information to follow and the hierarchical/reputation-based rules that may be introduced to further assist the process.

SQ-1.3: How well can the introduction of bid intercession enable a distributed decision-making system to centralise the decision-making process?

This last question investigates the potential resulting hybrid (centralised/distributed) behaviour of the system (given an asymmetric configuration/spread of resources). Being able to demonstrate the ability of the system to effectively emulate other decision-making architectures while being based on a distributed architecture would then open up avenues for fail-safe decision-making architectures with various degrees of centralisation. Furthermore, the performance of centralised and distributed MBA approaches is well understood and documented. As such, proving that the system is capable of effectively emulating a centralised system could allow for approximating its corresponding centralised performances through the transitive/equivalence property.

The corresponding hypothesis is therefore proposed below:

H-1.3 In a fully connected static network and a heterogeneous fleet with a single high-performance agent bidding for everyone with a higher priority, the method is capable of producing the same results as fully centralised methods (centred around a single high-performance agent).

This specific hypothesis was selected as it allows for verifying the fully centralised scenario, the other extreme decision-making distribution (fully distributed) being already guaranteed by the underlying consensus-based algorithm mechanism in conjunction with **H-1.1**.

This approach would also be investigated analytically and supported using a toy problem and scenario, specifically aimed at demonstrating the distributed/centralised behaviour (more in chapter 7).

SQ-2: How does the introduction of bid intercession impact task allocation performance?

This question addresses the need to understand the performance implications of being able to introduce further positive asymmetry in properties of the agents taking part in the decision-making process. This requires understanding the types and scale of PAs made possible, along with the performance implications of these. Two sub-questions are proposed:

SQ-2.1: What are the various positive asymmetries in agent properties enabled/impacted by bid intercession?

A large number of agent properties are constrained by on-the-edge requirements. The introduction of bid intercession allows for bypassing those, potentially introducing much larger capabilities in the decision-making process, and larger PAs in agent properties as a result. Establishing what those are is a critical first step in understanding the potential performance impact of bid intercession on a distributed decision-making process. This question therefore aims at investigating the various agent property PA enabled or impacted by intercession. Those notably include:

- Computational capability PA
- Information/cognitive PA
- Algorithm/bid evaluation method PA

The following driving hypothesis is hence proposed:

H-2.1 The introduction of bid intercession allows for greater positive asymmetry in a number of agent properties, which are not as easily achievable otherwise.

This question would be addressed theoretically.

SQ-2.2: How do the various agent property positive asymmetries impacted by bid intercession relate to task allocation performance?

Upon having established what agent property PAs are made possible by the introduction of bid intercession, understanding their impact on the decision-making process itself allows for better grasping of the potential decision-making process performance and quality implications. This question therefore focuses on understanding the relationship between the different PAs possible and their influence on the potential optimality and decision-making capabilities of a distributed network. The following driving hypothesis is proposed:

H-2.2 The scale of the various positive asymmetries in agent properties has a direct impact on the quality of the task allocation process. Larger positive asymmetries can in turn allow for a more effective decision-making process.

This would be again investigated using toy problems, further described in chapter 7.

SQ-3: How effective is bid intercession in providing external control over the decision-making process?

This question aims at investigating the potential of bid intercession in controlling the outcome of a distributed decision-making process. This is critical in helping to establish the potential use cases of such decision-making processes in a human-machine collaboration for example. Two main sub-questions were established:

SQ-3.1: How effective is bid intercession in injecting an external decision in the distributed decision-making process?

This question aims at investigating the ability of distributed decisions-processes to integrate external instructions. Given the nature of bid intercession, external inputs become potentially particularly easy to inject. Coupling that with the correct rule for which information source to prioritise could allow for an effective "piloting" of the decision-making process. This would prove critical in allowing for human input, a fundamental aspect of modern robotic systems. Demonstrating the ability to inject instructions in the decision-making process and having the fleet re-organise itself around that would enable advanced application and collaboration while retaining and enabling a significant level of autonomy. The corresponding hypothesis is therefore:

H-3.1 Given a fully connected network, the task-allocation process is capable of integrating an externally set task allocation within its broader plan and reorganising the task allocation around it effectively.

This hypothesis would be investigated experimentally using toy problems, also described in chapter 7.

7

Research Methodology

Having established the existing state-of-the-art methods and the research goals, it now becomes possible to establish a modelling approach. This chapter then focuses on the modelling techniques and tools adopted in section 7.1, along with the currently available tools used. More information is then provided on the experiments that shall be conducted and the data collection and subsequent analysis method adopted (section 7.3).

As discussed briefly in chapter 6, a number of experiment needs to be conducted to verify a number of hypotheses. The necessary experiments fall into two categories:

- **Toy problems:** Toy problems are heavily simplified, highly controlled scenarios, purposely designed to focus on gaining insight into specific aspects of a given problem. While not necessarily realistic, those experiments aim at investigating edge cases and boundaries of the solution space, in an attempt to effectively understand the dynamics of a given algorithm or approach.
- **Test scenarios:** Test scenarios aim to verify an approach's validity by running tests in a context as realistic as possible. Those allow for ensuring that the final systems would transfer well to real-world applications. To prove effective, the accent is put on creating a system which closely resembles the final one deployed, while substituting the environment and platforms it operates on using simulations.

To ensure consistent and reliable results are produced across all experiments, a modelling approach enabling both toy problems and test scenarios must be devised. A number of existing tools and frameworks can be leveraged to achieve this, with tools such as ROS2 for example ensuring broader compatibility with the rest of the robotics ecosystem and allowing for easy deployments on real robotics platforms. The modelling approach and architecture are therefore discussed in the next section.

7.1. Modelling Approach: ROS2 and MAF/CAF

To ensure an efficient and scalable modelling solution, the approach devised aims at building on top of existing and well-established tools and leveraging experience gathered from similar ventures previously. A common Agent-Based Modelling (ABM) approach methodology shall be adopted, which includes as main steps:

1. Definition of the environment
2. Definition of the agents
3. Definition of the agents' properties and internal processes
4. Definition of the agents' interactions. Two variants are considered here, namely CBAA and CBBA, with and without bid intercession. The intercession logic will also include the possibility of various hierarchical and reputation protocols.
5. Development of an event timeline generation. This will include the timestamp of each task, the location of the task, and the skill-set requirement for taking on said task (skills are derived from application/sensing capability).

6. Development of the various test environments used in toy problems (graph-based), and test scenarios (raster-based).
7. Development of the rest of the framework, including the ROS2 communication layer, the binding MBA framework, and the data collection, visualisation, and analysis.
8. Development of verification and validation strategies at both the local and global levels of the model.

To effectively simulate the various mechanisms at play, a time-step-based simulation is proposed. This has the double advantage of best representing reality, while also being easiest to port to non-simulated robotic platforms when necessary for further real-world testing. Additionally, the simulation framework is to be modular and flexible (avoiding having as many hard-coded elements as possible). This is done so to enable re-organising and adjusting simulation configurations as required by the various experimental setups. Consequentially, and to ensure the created framework scales well, the following requirements were established:

SIM- The simulation framework must operate in a time-step fashion The simulation framework must allow for any arbitrary agent dynamics (following a predefined template) with no further changes required to the rest of the simulation The simulation framework must allow for using any arbitrary consensus protocols (following a predefined template) with no further changes required to the rest of the simulation The simulation framework must allow for arbitrary fleet configurations with no further changes required to the rest of the simulation Arbitrary agent combination Arbitrary agent class The simulation framework must allow for using any arbitrary environments (following a predefined template) with no further changes required to the rest of the simulation The simulation framework must support dynamic environments

Based on the above-mentioned requirements and the availability of pre-existing tools and resources, a simulator architecture can be theorised. Four key components can be identified:

1. **The communication layer - ROS2:** ROS 2 is proposed as the main binding and communication layer. ROS 2, being the current standard in robotics, ensures a verified and validated foundation for the simulation and exchange of messages, and ensures compatibility of the created systems with other robotics systems and frameworks. It as such present itself as the natural choice for this application.

2. **Multi-agent Framework - MAF:** The multi-agent framework is the highest layer of abstraction in the simulation architecture and is responsible for the orchestration of the whole simulation, the exchange of messages, and the management of the interactions between the agents and the environment. It is an abstract software layer that may be used as a baseline by any application-specific experiments seeking to run multi-agent simulations. Abstracting away the base simulation layer has the advantage of allowing for the swapping of environments while keeping the ABM logic untouched, enabling a great level of re-usability.

3. **Common Auction Framework - CAF:** This layer inherits from the MAF layer to specifically focus on auction-based task allocation processes. This layer acts as a plug-and-play layer, allowing for various auction-based systems to be tested and hot-swapped, without needing to re-configure any other part of a given simulation setup.

4. **Thesis-specific modules:** Those modules are designed specifically for this research. They include the various environments used in the different experiments necessary to investigate all the hypotheses formulated. It also includes the different bid estimation logics, task allocation mechanisms (CBAA/CBBA), agents classes, and all other relevant blocks necessary for the various simulations.

A graphical representation of the above-described simulation architecture can be seen below in Figure 7.1. This architecture was devised and established based on a previous attempt made at the ONERA research to create a generalised and reusable MBA framework specifically for experimenting with auction-based processes. The first attempt was only partially successful but allowed for gaining insight into the various challenges faced when attempting to create such tools.

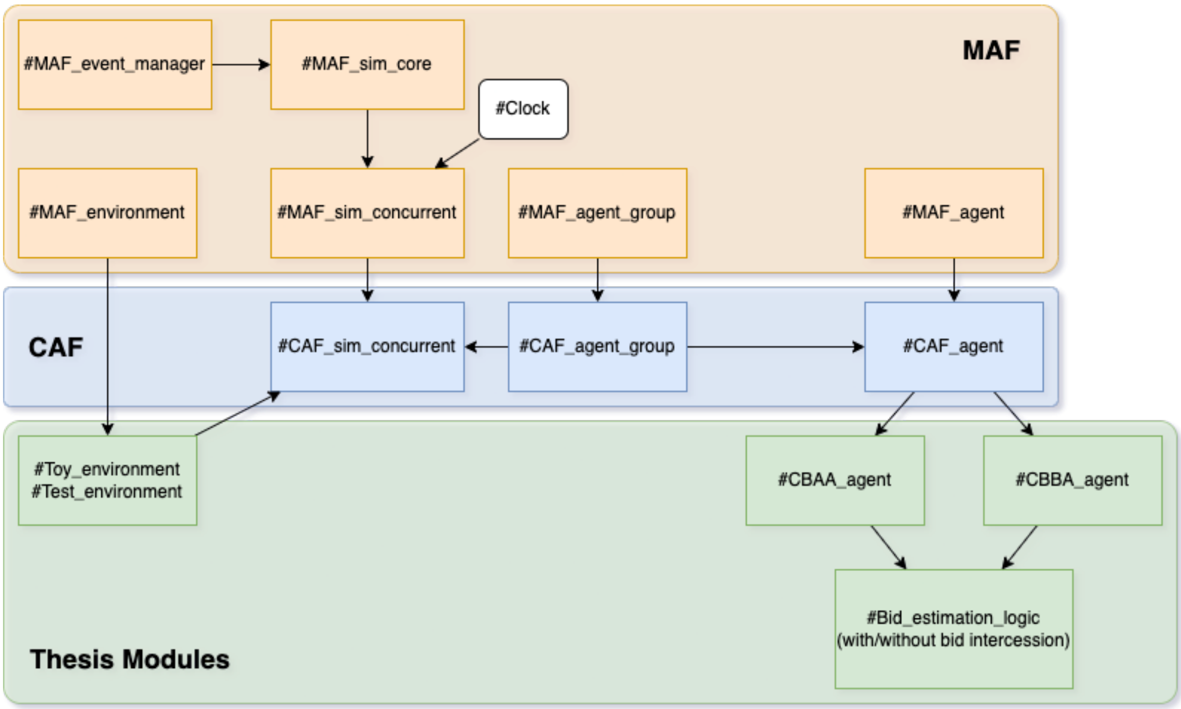

Figure 7.1: The three-layer architecture proposed for this research. The ROS@-based communication layer is not referenced here as it integrates across many modules.

7.2. Research Scope

To help frame the research well, the following areas of focus, simplifications, and assumptions are proposed:

- The study will restrain itself to static hierarchy systems. Demonstrating their stability and effectiveness would demonstrate the validity of the underlying principles, and in turn, open up the field to future research on dynamic hierarchy approaches and their interaction with consensus-based task allocation mechanisms.
- Uncertainty related to hardware will not be modelled. This notably includes sensor inaccuracies (sensors are assumed to always perform optimally)
- The tasks undertaken are assumed to be performed instantaneously once an agent with the correct skill-set reaches its location
- Dependencies constraints are ignored for this research. Each task is assumed to be independent of every other (ND scenarios)
- All agents are assumed to be cooperative. No agent will try to actively disrupt the network and act maliciously
- Agents will never get assigned a task they are not capable of undertaking
- Movement will be set to a fixed float value for the different agent classes, acceleration is not modelled
- Exploration is considered outside of the scope of this research. It is therefore assumed that exploration is performed by some other means, and the tasks are detected and issued to certain agents
- Communication is assumed to be instantaneous when available
- To obtain result sets allowing for verifying the various hypotheses investigated effectively, the experiments would have to be repeated for various environments and fleet configurations. Given the limited time and resources available for this research, the experimental data generated might have to be limited in scope. If this is the case, it will be specified in the results.

Additionally, all assumptions relevant to the CBAA and CBBA mechanisms are also maintained. This notably includes the Diminishing Marginal Gain assumption [11].

7.3. Data Collection and Analysis Methods

Upon having successfully set up the simulation framework discussed above, a number of different experiments will be conducted to investigate the various aspect of the research question. A Montecarlo method will be applied given the nature of the tests to ensure that the final results and conclusions drawn incorporate as much information and cover as many possible scenarios as possible.

As mentioned earlier, two categories of experiments will be performed. The first one, the toy problems, will aim at tackling specific aspects of the problem in an attempt to properly understand and frame the proposed algorithm's behaviours. The second (more complex) test will be to run the algorithm on a more realistic version of the search and rescue scenario, in an attempt to benchmark and compare the method's performance with other existing solutions. The two different experiment categories are discussed below.

7.3.1. Toy Problems

As discussed earlier in chapter 6, a number of hypotheses need to be investigated in order to fully address the various research questions. The hypotheses formulated are compiled below (the reference numbers match the research questions they correspond to):

H-1.1 The process of bid intercession does not have any impact on the convergence of the underlying distributed method **given that the source of truth is unambiguously defined at all times**. More specifically, the following are not impacted:
  - Ability to converge
  - Rate of convergence

H-1.3 In a fully connected static network and a heterogeneous fleet with a single high-performance agent bidding for everyone with a higher priority, the method is capable of producing the same results as fully centralised methods (centred around a single high-performance agent).

H-2.1 The introduction of bid intercession allows for greater positive asymmetry in a number of agent properties, which are not as easily achievable otherwise.

H-2.2 The scale of the various positive asymmetries in agent properties has a direct impact on the quality of the task allocation process. Larger positive asymmetries can in turn allow for a more effective decision-making process.

H-3.1 Given a fully connected network, the task-allocation process is capable of integrating an externally set task allocation within its broader plan and planning around it effectively.

Note that **H-2.1** on this list will be addressed theoretically, and as such does not require an experiment.

To best isolate the desired problem aspects, the experimental setup will be kept as simple as possible. A graph-based environment will be used as the basis of all tests, with different configurations and layouts as per the experiment's needs. The bid estimation process will also be kept as elementary as possible, with two main methods considered, aimed at representing two different cost estimation processes:

- **Manhattan Distance:** Bids are estimated on the basis of the inverse of the length of the shortest path between an agent and a given task
- **Weighted Manhattan Distance** Bids are estimated on the basis of the inverse of the length of the shortest path between an agent and a given task accounting for the weights of the edges.

While not significantly different, those two approaches allow for simulating a difference in knowledge, the first one only requiring information about the current position and task position, and the other also requiring knowledge of the graph itself. This in turn allows for introducing a difference in quality in the bid estimation process, which can in turn be leveraged to generate "conflicting" allocations.

Accordingly, four different toy experiments are proposed, each corresponding to a specific sub-research question hypothesis (the toy problem references matching the hypothesis addressed). A recap of the questions to hypothesis to toy problems references correspondences is provided below in Table 7.1.

Table 7.1: Overview of the research question to hypothesis to toy problem reference correspondence

Research question reference	**Sub-research question reference**	**Hypothesis reference**	**Toy problem reference**
3*SQ-1	SQ-1.1	H-1.1	T-1.1
	SQ-1.2	-	-
	SQ-1.3	H-1.3	T-1.3
2*SQ-2	SQ-2.1	H-2.1	-
	SQ-2.2	H-2.2	T-2.2
SQ-3	SQ-3.1	H-3.1	T-3.1

Toy problem T-1.1: Testing convergence

The goal of this first toy problem is to investigate the convergence and convergence rate properties of distributed consensus-based task allocation algorithms with intercession. As such, the experiment will be designed to specifically track the level and evolution of consensus within a network upon the discovery of a task by a given agent. The number of messages exchanged necessary to reach a consensus for various fleet and communication network configurations will be looked at in particular. Those tests will be run both for CBAA with and without intercession. Following the research scope limits set for this work, a hard hierarchy will be adopted here, and dynamic hierarchies will be left out of this study. If the hypothesis is correct, the results should indicate that no notable difference can be observed given that the consensus complexity (the number of agents participating in the consensus process) remains constant in both tests.

Toy problem T-1.3: Testing centralised emulation from distributed

The goal of the second toy problem is to investigate the ability of a distributed network to operate in a centralised fashion. Accordingly, the hierarchy within the network will be designed to reflect a centralised system, with all agents but one place at the same priority level, and a single one being placed higher. No constraints will be placed on the communication network, meaning all agents will be able to exchange messages freely and reliably. The aim of the experiment will then be to verify whether, in such conditions, the network converges towards the solutions set by the higher priority agent consistently. If this hypothesis is correct (and given the scenario defined in the hypothesis), the experiments should demonstrate that all bids considered in the auctions were emitted by the central agent responsible for the centralised coordination.

Toy problem T-2.2: Testing positive asymmetries

This experiment is the most complex one to design, as it requires coming up with experimental setups reflecting all the positive asymmetries investigated. The exact configuration of each test will therefore not only depend on the results of **SQ-2.1**, which will dictate what PA are worth investigating, but also on the technical and computational capabilities available for this research. The general approach however will be as follows. A reference experiment will be set up and run to obtain a baseline to compare other results to. This base experimental setup will then be maintained, but the various agents' capabilities will be adjusted to reflect the desired PA respectively. The results will then be collected and compared to determine the relative impact of the various PA on the system's performance.

Toy problem T-3.1: Testing external intervention capabilities

The goal of this final toy problem is to investigate the ability of the system to accept external instructions, and seamlessly integrate them within its own plan. To demonstrate this, a specific agent/task combination will be selected, and a bid will be injected mid-run. The goal will then be to verify the instruction is respected, and how well the network re-organises itself around it. It should be noted that this experiment, while seemingly similar to T-1.2, diverges in the fact that only a single instruction is injected, whereas in T-1.2 an entire plan is imposed.

7.3.2. Search and Rescue case study

Upon having performed and investigated all the toy problems described earlier, a final test need to be run, with the goal of evaluating the algorithm's performance within a more realistic context. Following the simulation methodology described in chapter 2, the following experimental setup is proposed.

Environment

The environment proposed for this case study is reconstructed based on a satellite image of one of ONERA's test sites (Figure 7.2). This particular area was selected for a number of reasons. First, the site presents a large variety of features relevant to the desired scenario. Those include a blend of dense and light vegetation, which can be used to help stimulate communication. Furthermore, this site belonging to the ONERA could eventually allow for performing on-site experiments, which would in turn be used to validate the results obtained in this research.

Figure 7.2: The Caylus military field, used as reference scenario for this research

The goal of this feature collection was to be able to distinguish between various map features and provide each category with a unique set of characteristics. This process was done manually and produced a set of svg masks. A total of four different feature sets were extracted:

1. **Hard obstacles:** Those notably include all buildings, hard infrastructure, and containers present on the site
2. **Dense vegetation:** Those include the trees and all dense forest areas

3. **Light vegetation:** Those include the small bushes, and tall grass areas
4. **Paths:** Those includes all roads and dirt paths indiscriminately

These were aggregated and are displayed below in Figure 7.3 according to the list order above from darkest to lightest:

Figure 7.3: The final feature maps extracted from the Caylus satellite image agregated

The map created will therefore allow for simulating the movement of agents on the map, and the communication constraints experienced (each obstacle category being provided with different levels of communication permeability).

Agents

The agents shall be simulated as described in chapter 2, and will endorse the responsibilities as laid out in the RoboCupRescue challenge.

Coordination protocols

A total of four coordination protocols will be tested. Those include CBAA and CBBA without intercession, to obtain baseline performances to use as a point of comparison, and CBAA and CBBA with intercession.

Coordination architecture

Upon having decided on the coordination protocols, it is necessary to define the overall communication and coordination structure adopted between the various agent categories. This notably includes which agents will take part in which census processes, which agents are responsible for which task categories, and ultimately which tasks to analyse during the consensus process. Although not absolutely necessary for the correct functioning of the underlying algorithms, this ensures agents only perform computations on the subset of tasks they need for their own work, reducing the computational load where possible. In such architectures, all agents still relay all tasks and bids but only perform the consensus process on the tasks relevant to them. An example of such optimisation would be to have the "pick up a victim" tasks and bids only be evaluated by the ambulance agents, and have them in turn not take part in the consensus on the "go extinguish a fire" tasks, which are only relevant to the firefighter agents. For the later tasks and bids, the ambulance agents simply act as proxies, ensuring the effective flow of information without interfering in the consensus process.

This allows for maintaining a robust and fully-populated communication network leveraging all possible agents to support the flow of information while only performing the necessary calculations where relevant. Additionally, this results in the consensus process only being performed among relevant subgroups of agents, in turn confining each consensus processes complexities to the minimum possible (it being directly related to the size of each relevant agent subgroup). This in turn allows for faster convergence to a decision as fewer agents take part in it.

Finally, the structuring of the various consensus processes allows for controlling the coordination process and assigning responsibility to various agent subgroups, further structuring the coordination process as needed. For this specific research, and given that it was stated that none of the tasks had any dependencies on each other, it is sufficient to simply constrain each task and bid categories to their relevant agent subgroups, as no additional structure is needed. It should be noted that if the no-dependency assumption was not made, it could be possible to further optimise the network architecture, by for example introducing "coordination hierarchies". For instance, we could have each agent group "centres" (ambulance centre, police station, fire station) communicating among themselves and developing strategies for addressing the above-mentioned dependencies in the task allocation before passing those on to the agent network through intercession for example, leveraging the centralisation capability of the system for the decision-making process.

Explain the intercession architecture

8 Research Planning

This chapter provides an overview of the logistics for this research project. A breakdown of the timeline and major milestones is first provided in section 8.1. The resource and budget allocation is then presented in section 8.2, and finally, the potential risks and corresponding mitigation strategies are discussed in section 8.3.

8.1. Timeline and Milestones for the Research Project

To help structure and guide the work undertaken for this research, five major project phases have been established:

1. **Literature Study:** The first phase of the research project is the literature study. Its main purpose is to provide an overview of the current state of the art, identify a research gap, research questions, and a corresponding research methodology. The final result of this phase of the project is the literature review paper (this document).

2. **Research and Development:** The research and development phase will be done following the conclusions and methodology established during the literature study. The main goals of this phase will be to setup all the digital infrastructure and frameworks necessary for running the experiments during the testing phase. It should be noted that no particular date was planned for the various sub-components of the development process, it was decided instead to follow an agile methodology for this part of the work.

3. **Testing:** Upon having completed all necessary tools and simulation frameworks, the experiments defined in the literature study will be conducted, and data will be collected for subsequent analysis.

   This phase is notably broken down into two major sub-phases:

   (a) Toy problems: A number of toy problems will first be researched, verifying and validating a number of specific hypotheses made on various aspects of the research problem

   (b) Case study: A main experiment, focusing on testing the novel method on the SAR problem will then be conducted

4. **Data Analysis:** The data collected during the testing phase will then be analysed, and used to attempt at verifying the hypotheses made to address the main research question.

5. **Reporting:** The final phase of this project will be to report all results, findings, and conclusions in a final thesis paper.

A breakdown table of the above phases is provided bellow in Table 8.1:

Table 8.1: Breakdown of the research phases

Project Phases

Task	Timeline	Duration	Status	Label	Dependent On
Literature Study phase 4	Feb 13 - May 12	89	Ongoing	Project p...	-

Subitem	Status	Label	Dependent On
Screening of papers	Pending	Task	-
Paper classification and ...	Pending	Task	Screening of papers
Definition of Research Q...	Ongoing	Task	-
Research Methodologies	Ongoing	Task	-

Task	Timeline	Duration	Status	Label	Dependent On
Design and Development phase 5	Mar 15 - Jun 6	84	Ongoing	Project p...	Literature Study phase

Subitem	Timeline	Duration	Status	Label	Dependent On
Develop agents ABC	-		Pending	Task	-
Develop action agents	-		Pending	Task	Develop agents ABC
Develop coordination ag...	-		Pending	Task	Develop agents ABC
Develop scenario generat...	-		Pending	Task	-
Develop coordination-lay...	-		Pending	Task	Develop agents ABC

Task	Timeline	Duration	Status	Label	Dependent On
Testing phase 2	Jun 10 - Jul 26	47	Pending	Project p...	Design and Development ph...

Subitem	Timeline	Duration	Status	Label	Dependent On
Toy problems	-		Pending	Task	-
Scenario run	-		Pending	Task	-

Task	Timeline	Duration	Status	Label	Dependent On
Data Analysis phase 3	Jul 26 - Aug 25	31	Pending	Project p...	Testing phase

Subitem	Timeline	Duration	Status	Label	Dependent On
Cleanup/formating	-		Pending	Task	-
Post-processing	-		Pending	Task	Cleanup/formating
Synthesis/statistical anal...	-		Pending	Task	Post-processing

Accompanying the project phases, a number of key milestones have been defined for this research. Their main goal is to serve as yardsticks, aimed at evaluating the progress of the project as the research is conducted. Those milestones are:

1. **Kick-off meeting:** The kick-off meeting is the first major milestone for this project. The goal of the kick-off meeting will be to perform an evaluation of the literature study performed. The meeting will notably focus on the research proposal, the research methodology developed around it, and the questions and objectives of the project. It will officially mark the start of the thesis project. Additionally, tentative goals for the mid-term meeting will be set.

2. **Literature study submission:** The submission of the final literature study paper is the next major milestone. The submission should take place roughly one week after the kick-off meeting, which will be used to update the first draft of the paper to account for the comments and conclusions reached during the meeting.

3. **Mid-term meeting:** The mid-term meeting will occur three months after the kick-off meeting. A presentation will be given, summarising the approach, and providing details on the methodological steps and results obtained so far. Additionally, the next steps of the work will be identified in detail to ensure the successful completion of the project. A small report summarising the work

achieved so far will also be provided here.

4. **Thesis first draft:** The completion of the first draft of the thesis will then be the next major goal. This is necessary for the initiation of the green light review. It will allow for obtaining feedback on the work and results obtained, and allow for starting the final correction process.
5. **Green light review:** The purpose of the green light review will be to review the state of the research and evaluate whether the work achieved so far is sufficient for starting the submission and review process.
6. **Thesis submission:** The final draft of the thesis, taking into account all the feedback received will then be submitted for review and defence.
7. **Defence:** The thesis defence is the final step of the process. The research will be presented, and a board of examiners will be given the opportunity to ask all necessary questions and doubts before performing a final evaluation of the overall project.

Table 8.2: Recap of the research project milestones

Milestones

Task	Timeline	Duration	Status	Label	Dependent On
Kick-off meeting	May 19	0	Pending	Milestone	-
Literature Study submission	May 25	0	Ongoing	Milestone	Literature Study phase
Research Methodologies submission	Jun 30	0	Ongoing	Milestone	-
Mid-term meeting	Aug 14	0	Pending	Milestone	Literature Study submission
Thesis draft	Sep 10	0	Pending	Milestone	Mid-term meeting
Green light review (go/no go)	Oct 15	0	Pending	Milestone	Thesis draft
Examination request	Oct 20	0	Pending	Milestone	Green light review (go/no go)
Thesis submission	Oct 25	0	Pending	Milestone	Examination request
Defense	Nov 1	0	Pending	Milestone	Examination request
Reporting phase	Jun 21 - Oct 23	125	Pending	Project p...	Design and Development ph...
		376 sum			

Schedule

Task	Timeline	Duration	Status	Label	Dependent On
ONERA access	Feb 13 - Aug 4	173	Ongoing	Event	-
		173 sum			

A Gantt chart recapping the phases and milestones can be seen below.

Thesis

Read-only view, generated on 20 Jul 2023

Instagantt

	ACTIVITIES	START	CD	DUE	DPD
	Milestones:	10/Jul	134d	20/Nov	
1	Kick-off meeting	10/Jul	1d	10/Jul	11
2	Literature Study submission	21/Jul	1d	21/Jul	1, 11
3	Mid-term meeting	14/Aug	1d	14/Aug	2
4	Research Methodologies su...	24/Aug	1d	24/Aug	3
5	Thesis draft	04/Oct	1d	04/Oct	4
6	Green light review (go/no go)	15/Oct	1d	15/Oct	5
7	Examination request	20/Oct	1d	20/Oct	6
8	Thesis submission	25/Oct	1d	25/Oct	7, 32
9	Defense	20/Nov	1d	20/Nov	8
	Project Phases:	13/Feb	253d	23/Oct	
	Literature Study phase	13/Feb	138d	30/Jun	
12	Screening of papers		1d		
13	Paper classification and...		1d		12
14	Definition of Research ...		1d		13
15	Research Methodologies		1d		14
	Design and Development p...	15/Mar	139d	31/Jul	
17	Develop agents ABC		1d		
18	Develop agents		1d		
19	Develop communicatio...		1d		
20	Develop environments		1d		
21	Develop scenario gener...		1d		
22	Develop MonteCarlo ru...		1d		21
	Testing phase	02/Aug	30d	31/Aug	
	Toy problems	02/Aug	26d	27/Aug	
25	H-1.1	02/Aug	4d	05/Aug	
26	H-1.3	06/Aug	6d	11/Aug	
27	H-2.1	12/Aug	6d	17/Aug	
28	H-2.2	18/Aug	8d	25/Aug	
29	H-3.1	26/Aug	2d	27/Aug	
30	Scenario run	28/Aug	4d	31/Aug	
31	Data Analysis phase	01/Sep	30d	30/Sep	
32	Reporting phase	01/Sep	53d	23/Oct	

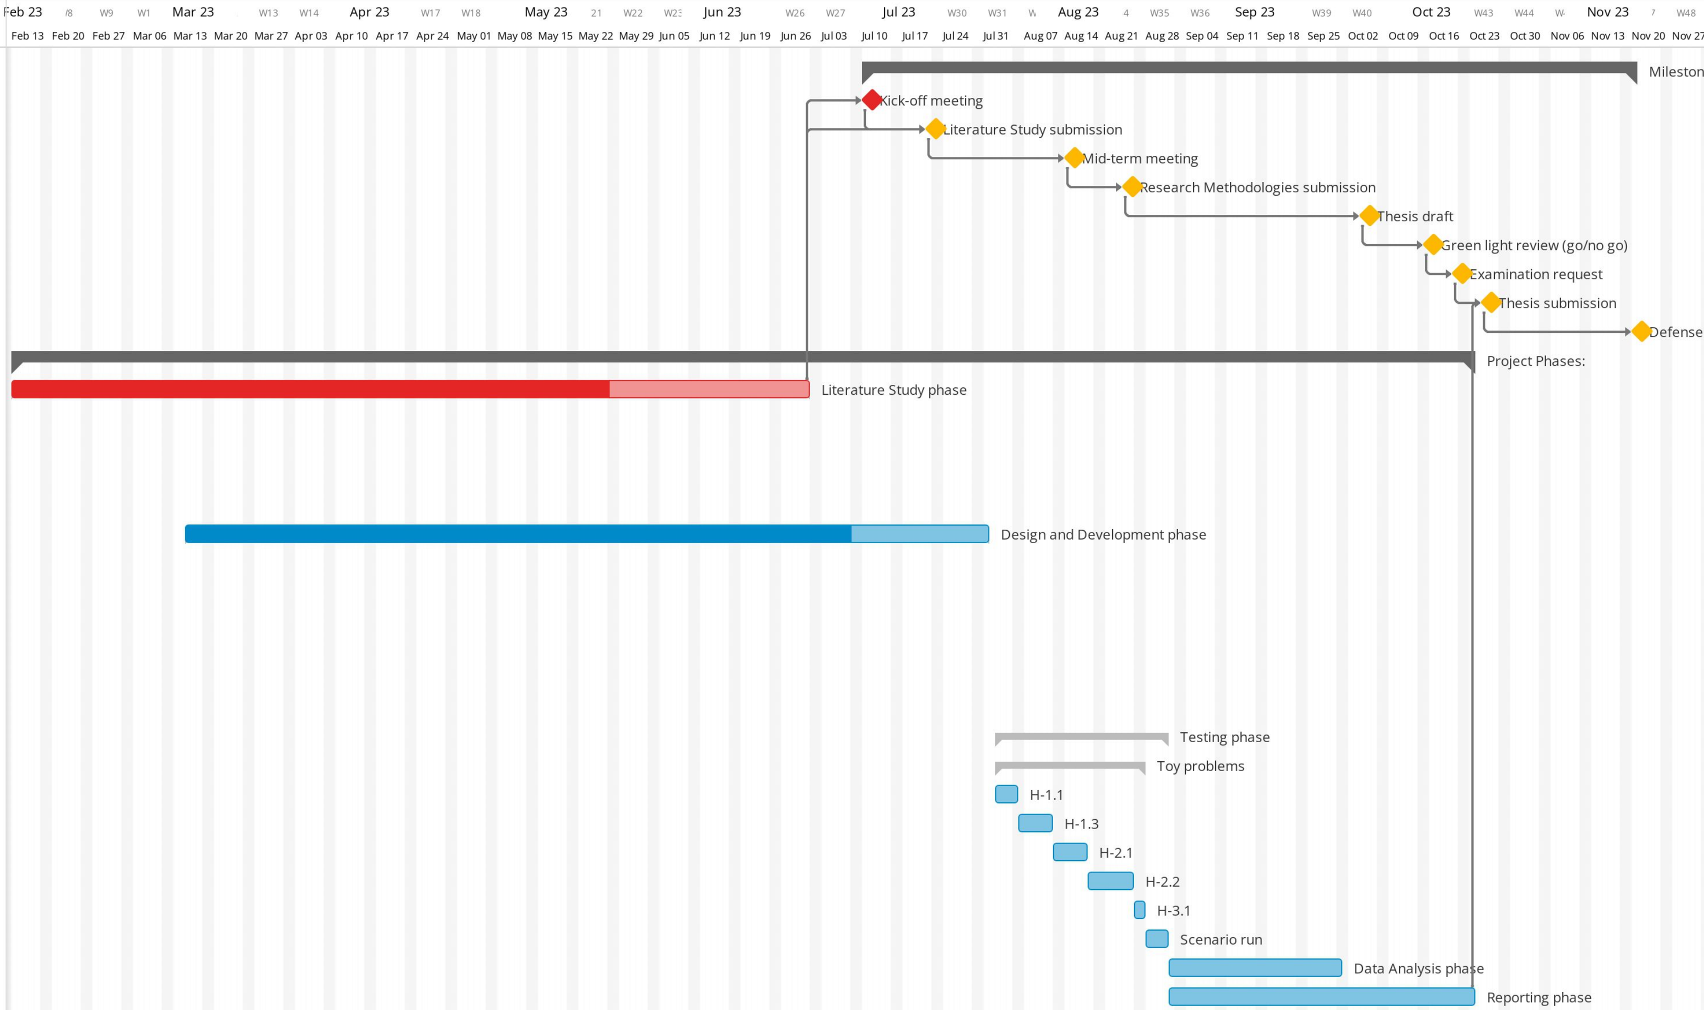

8.2. Resource Allocation and Budget

The research project is performed in cooperation with the Office National d'Etude et de Recherche Aerospatial (ONERA, a.k.a the French Aerospace Labs) research centre. They will provide the necessary computing resources along with financing for a duration of six months. If the research timeline allows for it, the ONERA may furthermore provide the opportunity to conduct a number of real-life tests in their Voliere multi-agent research lab.

8.3. Potential Risks and Mitigation Strategies

The main risk is the potential for delays related to the development cycle of the tools and framework necessary for conducting the various experiments. This was however accounted for by performing a large majority of the baseline development while writing the literature study. This in turn not only allowed to obtain a number of preliminary confirmations with respect to the topic itself but also allowed for obtaining a good understanding of what would be realistic given the limited timeline and resources.

Bibliography

[1] Boston Dynamics Partners With OTTO Motors To Coordinate Mobile Robots In The Warehouse, . URL https://www.bostondynamics.com/press-release-boston-dynamics-2020-03-03.

[2] Documentation – RoboCup Rescue Simulation, . URL https://rescuesim.robocup.org/resources/documentation/.

[3] Representing reputation in a multiagent system – LINCS, . URL https://www.lincs.fr/events/representing-reputation-in-a-multiagent-system/.

[4] RoboCup Federation official website, . URL https://www.robocup.org/.

[5] Dirichlet distribution, April 2023. URL https://en.wikipedia.org/w/index.php?title=Dirichlet_distribution&oldid=1149273862. Page Version ID: 1149273862.

[6] The Future of Robotics: Orchestrating the Heterogeneous Robot Fleet, May 2023. URL https://blogs.gartner.com/power-of-the-profession-blog/the-future-of-robotics-orchestrating-the-heterogeneous-robot-fleet/.

[7] Guruprasad Airy, Tracy Mullen, and John Yen. Market based adaptive resource allocation for distributed rescue teams. April 2009.

[8] Matthew Argyle, David W. Casbeer, and Randy Beard. A Multi-Team Extension of the Consensus-Based Bundle Algorithm. In Proceedings of the 2011 American Control Conference, pages 5376–5381, San Francisco, CA, June 2011. IEEE. ISBN 978-1-4577-0081-1 978-1-4577-0080-4 978-1-4577-0079-8. doi: 10.1109/ACC.2011.5991162. URL http://ieeexplore.ieee.org/document/5991162/.

[9] Salvatore Aronica, Francesco Benvegna, Massimo Cossentino, Salvatore Gaglio, Alessio Langiu, Carmelo Lodato, Salvatore Lopes, Umberto Maniscalco, and Pierluca Sangiorgi. An Agent-based System for Maritime Search and Rescue Operations. volume 621, September 2010.

[10] Amir Jalaly Bidgoly and Behrouz Tork Ladani. Quantitative verification of beta reputation system using PRISM probabilistic model checker. In 2013 10th International ISC Conference on Information Security and Cryptology (ISCISC), pages 1–6, August 2013. doi: 10.1109/ISCISC.2013.6767336.

[11] Luc Brunet, Han-Lim Choi, and Jonathan How. Consensus-Based Auction Approaches for Decentralized Task Assignment. In AIAA Guidance, Navigation and Control Conference and Exhibit, Honolulu, Hawaii, August 2008. American Institute of Aeronautics and Astronautics. ISBN 978-1-60086-999-0. doi: 10.2514/6.2008-6839. URL https://arc.aiaa.org/doi/10.2514/6.2008-6839.

[12] Lei Cao, He shun Tan, Hui Peng, and Ming cong Pan. Multiple UAVs hierarchical dynamic task allocation based on PSO-FSA and decentralized auction. In 2014 IEEE International Conference on Robotics and Biomimetics (ROBIO 2014), pages 2368–2373, December 2014. doi: 10.1109/ROBIO.2014.7090692.

[13] Han-Lim Choi, Luc Brunet, and Jonathan P. How. Consensus-Based Decentralized Auctions for Robust Task Allocation. IEEE Transactions on Robotics, 25(4):912–926, August 2009. ISSN 1941-0468. doi: 10.1109/TRO.2009.2022423. Conference Name: IEEE Transactions on Robotics.

[14] José Delgado-Penín. Radio Communication Systems simulation: from the pioneers to the present. September 2015. doi: 10.1109/EUROCON.2015.7313728.

[15] Luis Emmi, Mariano Gonzalez-de Soto, Gonzalo Pajares, and Pablo Gonzalez-de Santos. New Trends in Robotics for Agriculture: Integration and Assessment of a Real Fleet of Robots. The Scientific World Journal, 2014:e404059, March 2014. ISSN 2356-6140. doi: 10.1155/2014/404059. URL https://www.hindawi.com/journals/tswj/2014/404059/. Publisher: Hindawi.

[16] Saurabh Ganeriwal, Laura K. Balzano, and Mani B. Srivastava. Reputation-based framework for high integrity sensor networks. ACM Transactions on Sensor Networks, 4(3):15:1–15:37, June 2008. ISSN 1550-4859. doi: 10.1145/1362542.1362546. URL https://dl.acm.org/doi/10.1145/1362542.1362546.

[17] Nikhil Garg and Ramesh Johari. Designing Optimal Binary Rating Systems. In Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics, pages 1930–1939. PMLR, April 2019. URL https://proceedings.mlr.press/v89/garg19a.html. ISSN: 2640-3498.

[18] Brian P. Gerkey and Maja J. Matarić. A Formal Analysis and Taxonomy of Task Allocation in Multi-Robot Systems. The International Journal of Robotics Research, 23(9):939–954, September 2004. ISSN 0278-3649, 1741-3176. doi: 10.1177/0278364904045564. URL http://journals.sagepub.com/doi/10.1177/0278364904045564.

[19] Jones Granatyr, Vanderson Botelho, Otto Lessing, Edson Scalabrin, Jean-Paul Barthès, and Fabrício Enembreck. Trust and Reputation Models for Multi-Agent Systems. ACM Computing Surveys, 48, October 2015. doi: 10.1145/2816826.

[20] Navid Hooshangi, Ali Asghar Alesheikh, Mahdi Panahi, and Saro Lee. Urban search and rescue (USAR) simulation system: spatial strategies for agent task allocation under uncertain conditions. Natural Hazards and Earth System Sciences, 21(11):3449–3463, November 2021. ISSN 1561-8633. doi: 10.5194/nhess-21-3449-2021. URL https://nhess.copernicus.org/articles/21/3449/2021/. Publisher: Copernicus GmbH.

[21] Simon Hunt, Qinggang Meng, Chris Hinde, and Tingwen Huang. A Consensus-Based Grouping Algorithm for Multi-agent Cooperative Task Allocation with Complex Requirements. Cognitive Computation, 6(3):338–350, September 2014. ISSN 1866-9956, 1866-9964. doi: 10.1007/s12559-014-9265-0. URL http://link.springer.com/10.1007/s12559-014-9265-0.

[22] Ahmed Hussein, Mohamed Abdelhady, Mohamed Bakr, Omar Shehata, and Alaa Khamis. Multi-robot Task Allocation for Search and Rescue Missions. Journal of Physics: Conference Series, 570: 1–10, December 2014. doi: 10.1088/1742-6596/570/5/052006.

[23] Roslan Ismail and Audun Josang. The Beta Reputation System.

[24] Sarah Ismail and Liang Sun. Decentralized hungarian-based approach for fast and scalable task allocation. In 2017 International Conference on Unmanned Aircraft Systems (ICUAS), pages 23–28, June 2017. doi: 10.1109/ICUAS.2017.7991447.

[25] Luke Johnson, Sameera Ponda, Han-lim Choi, and Jonathan How. Improving the Efficiency of a Decentralized Tasking Algorithm for UAV Teams with Asynchronous Communications. In AIAA Guidance, Navigation, and Control Conference, Guidance, Navigation, and Control and Co-located Conferences. American Institute of Aeronautics and Astronautics, August 2010. doi: 10.2514/6.2010-8421. URL https://arc-aiaa-org.tudelft.idm.oclc.org/doi/10.2514/6.2010-8421.

[26] Luke Johnson, Sameera Ponda, Han-Lim Choi, and Jonathan How. Asynchronous Decentralized Task Allocation for Dynamic Environments. In Infotech@Aerospace 2011, St. Louis, Missouri, March 2011. American Institute of Aeronautics and Astronautics. ISBN 978-1-60086-944-0. doi: 10.2514/6.2011-1441. URL https://arc.aiaa.org/doi/10.2514/6.2011-1441.

[27] Audun Josang and Jochen Haller. Dirichlet Reputation Systems. In The Second International Conference on Availability, Reliability and Security (ARES'07), pages 112–119, Vienna, Austria, 2007. IEEE. ISBN 978-0-7695-2775-8. doi: 10.1109/ARES.2007.71. URL http://ieeexplore.ieee.org/document/4159794/.

[28] Audun Jøsang. Subjective Logic. Artificial Intelligence: Foundations, Theory, and Algorithms. Springer International Publishing, Cham, 2016. ISBN 978-3-319-42335-7 978-3-319-42337-1. doi: 10.1007/978-3-319-42337-1. URL http://link.springer.com/10.1007/978-3-319-42337-1.

[29] Audun Jøsang, Ross Hayward, and Simon Pope. Trust Network Analysis with Subjective Logic.

[30] Cihan Kaleli and Huseyin Polat. Robustness Analysis of Naïve Bayesian Classifier-Based Collaborative Filtering. In Christian Huemer and Pasquale Lops, editors, E-Commerce and Web Technologies, Lecture Notes in Business Information Processing, pages 202–209, Berlin, Heidelberg, 2013. Springer. ISBN 978-3-642-39878-0. doi: 10.1007/978-3-642-39878-0_19.

[31] Sven Koenig. The Power of Sequential Single-Item Auctions for Agent Coordination. page 5.

[32] G. Ayorkor Korsah, Anthony Stentz, and M. Bernardine Dias. A comprehensive taxonomy for multi-robot task allocation. The International Journal of Robotics Research, 32(12):1495–1512, October 2013. ISSN 0278-3649, 1741-3176. doi: 10.1177/0278364913496484. URL http://journals.sagepub.com/doi/10.1177/0278364913496484.

[33] Nathan Lindsay, Russell K. Buehling, and Liang Sun. A Sequential Task Addition Distributed Assignment Algorithm for Multi-Robot Systems. Journal of Intelligent & Robotic Systems, 102 (2):51, June 2021. ISSN 0921-0296, 1573-0409. doi: 10.1007/s10846-021-01394-2. URL https://link.springer.com/10.1007/s10846-021-01394-2.

[34] Yining Liu, Keqiu Li, Yingwei Jin, Yong Zhang, and Wenyu Qu. A novel reputation computation model based on subjective logic for mobile ad hoc networks. Future Generation Computer Systems, 27(5):547–554, May 2011. ISSN 0167-739X. doi: 10.1016/j.future.2010.03.006. URL https://www.sciencedirect.com/science/article/pii/S0167739X10000518.

[35] Bernard Marr. The Best Examples Of Human And Robot Collaboration. Forbes. URL https://www.forbes.com/sites/bernardmarr/2022/08/10/the-best-examples-of-human-and-robot-collaboration/. Section: Enterprise Tech.

[36] Pranab Muhuri and Amit Rauniyar. Immigrants Based Adaptive Genetic Algorithms for Task Allocation in Multi-Robot Systems. International Journal of Computational Intelligence and Applications, 16:1750025, December 2017. doi: 10.1142/S1469026817500250.

[37] L. Mui, Mojdeh Mohtashemi, and A. Halberstadt. A computational model of trust and reputation. pages 2431–2439, February 2002. doi: 10.1109/HICSS.2002.994181.

[38] Antonio Neves. Service Robots. BoD – Books on Demand, January 2018. ISBN 978-953-51-3722-1. Google-Books-ID: I8WPDwAAQBAJ.

[39] James Parker, Ernesto Nunes, Julio Godoy, and Maria Gini. Exploiting Spatial Locality and Heterogeneity of Agents for Search and Rescue Teamwork*. Journal of Field Robotics, 33(7):877–900, 2016. ISSN 1556-4967. doi: 10.1002/rob.21601. URL http://onlinelibrary.wiley.com/doi/abs/10.1002/rob.21601. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/rob.21601.

[40] Ruchir Patel, Eliot Rudnick-Cohen, Shapour Azarm, Michael Otte, Huan Xu, and Jeffrey W. Herrmann. Decentralized Task Allocation in Multi-Agent Systems Using a Decentralized Genetic Algorithm. In 2020 IEEE International Conference on Robotics and Automation (ICRA), pages 3770–3776, Paris, France, May 2020. IEEE. ISBN 978-1-72817-395-5. doi: 10.1109/ICRA40945.2020.9197314. URL https://ieeexplore.ieee.org/document/9197314/.

[41] Gauthier Picard. Multi-Agent Consensus-based Bundle Allocation for Multi-Mode Composite Tasks. May 2023. URL https://hal.science/hal-03964386.

[42] Sameera Ponda, Olivier Huber, Han-Lim Choi, and Jonathan P. How. Avoid communication outages in decentralized planning. In 2010 IEEE Globecom Workshops, pages 1756–1759, December 2010. doi: 10.1109/GLOCOMW.2010.5700242. ISSN: 2166-0077.

[43] Sameera Ponda, Nisar Ahmed, Brandon Luders, Eric Sample, Tauhira Hoossainy, Danelle Shah, Mark Campbell, and Jonathan How. Decentralized Information-Rich Planning and Hybrid Sensor Fusion for Uncertainty Reduction in Human-Robot Missions. In AIAA Guidance, Navigation, and Control Conference, Portland, Oregon, August 2011. American Institute of Aeronautics and Astronautics. ISBN 978-1-60086-952-5. doi: 10.2514/6.2011-6238. URL https://arc.aiaa.org/doi/10.2514/6.2011-6238.

[44] Jorge Peña Queralta, Jussi Taipalmaa, Bilge Can Pullinen, Victor Kathan Sarker, Tuan Nguyen Gia, Hannu Tenhunen, Moncef Gabbouj, Jenni Raitoharju, and Tomi Westerlund. Collaborative Multi-Robot Systems for Search and Rescue: Coordination and Perception, August 2020. URL http://arxiv.org/abs/2008.12610. arXiv:2008.12610 [cs].

[45] Felix Quinton, Christophe Grand, and Charles Lesire. Market Approaches to the Multi-Robot Task Allocation Problem: a Systematic Mapping and Survey. page 48.

[46] A.A. Selcuk, E. Uzun, and M.R. Pariente. A reputation-based trust management system for P2P networks. In IEEE International Symposium on Cluster Computing and the Grid, 2004. CCGrid 2004., pages 251–258, April 2004. doi: 10.1109/CCGrid.2004.1336575.

[47] Mark Selden, Jason Zhou, Felipe Campos, Nathan Lambert, Daniel Drew, and Kristofer S. J. Pister. BotNet: A Simulator for Studying the Effects of Accurate Communication Models on Multi-agent and Swarm Control, August 2021. URL http://arxiv.org/abs/2108.13606. arXiv:2108.13606 [cs].

[48] Boris Skoric, Sebastiaan J. A. de Hoogh, and Nicola Zannone. Flow-based reputation with uncertainty: Evidence-Based Subjective Logic, February 2015. URL http://arxiv.org/abs/1402.3319. arXiv:1402.3319 [cs].

[49] P.B. Sujit and Randy Beard. Distributed Sequential Auctions for Multiple UAV Task Allocation. In 2007 American Control Conference, pages 3955–3960, July 2007. doi: 10.1109/ACC.2007.4282558. ISSN: 2378-5861.

[50] Jian Tang, Kejun Zhu, Haixiang Guo, Can Liao, and Shuwen Zhang. Simulation Optimization of Search and Rescue in Disaster Relief Based on Distributed Auction Mechanism. Algorithms, 10 (4):125, December 2017. ISSN 1999-4893. doi: 10.3390/a10040125. URL https://www.mdpi.com/1999-4893/10/4/125. Number: 4 Publisher: Multidisciplinary Digital Publishing Institute.

[51] Sahar Trigui, Anis Koubaa, Omar Cheikhrouhou, Habib Youssef, Hachemi Bennaceur, Mohamed-Foued Sriti, and Yasir Javed. A Distributed Market-based Algorithm for the Multi-robot Assignment Problem. Procedia Computer Science, 32:1108–1114, January 2014. ISSN 1877-0509. doi: 10.1016/j.procs.2014.05.540. URL https://www.sciencedirect.com/science/article/pii/S1877050914007406.

[52] Antidio Viguria, Ivan Maza, and Anibal Ollero. SET: An algorithm for distributed multirobot task allocation with dynamic negotiation based on task subsets. In Proceedings 2007 IEEE International Conference on Robotics and Automation, pages 3339–3344, April 2007. doi: 10.1109/ROBOT.2007.363988. ISSN: 1050-4729.

[53] A. Vosoughi, J. R. Cavallaro, and A. J. Marshall. Trust-aware Consensus-inspired Distributed Cooperative Spectrum Sensing for Cognitive Radio Ad Hoc Networks. IEEE Transactions on Cognitive Communications and Networking, 2(1):24–37, June 2016. URL https://ieeexplore.ieee.org/document/7497578. Number: 1 Publisher: Institute of Electrical and Electronics Engineers (IEEE).

[54] Amanda Whitbrook, Qinggang Meng, and Paul Chung. A Robust, Distributed Task Allocation Algorithm for Time-Critical, Multi Agent Systems Operating in Uncertain Environments. June 2017. ISBN 978-3-319-60044-4. doi: 10.1007/978-3-319-60045-1_8.

[55] Andrew K. Whitten, Han-Lim Choi, Luke B. Johnson, and Jonathan P. How. Decentralized Task Allocation with Coupled Constraints in Complex Missions. In Proceedings of the 2011

American Control Conference, pages 1642–1649, San Francisco, CA, June 2011. IEEE. ISBN 978-1-4577-0081-1 978-1-4577-0080-4 978-1-4577-0079-8. doi: 10.1109/ACC.2011.5990917. URL http://ieeexplore.ieee.org/document/5990917/.

[56] Wanqing Zhao, Qinggang Meng, and Paul W. H. Chung. A Heuristic Distributed Task Allocation Method for Multivehicle Multitask Problems and Its Application to Search and Rescue Scenario. IEEE Transactions on Cybernetics, 46(4):902–915, April 2016. ISSN 2168-2275. doi: 10.1109/TCYB.2015.2418052. Conference Name: IEEE Transactions on Cybernetics.