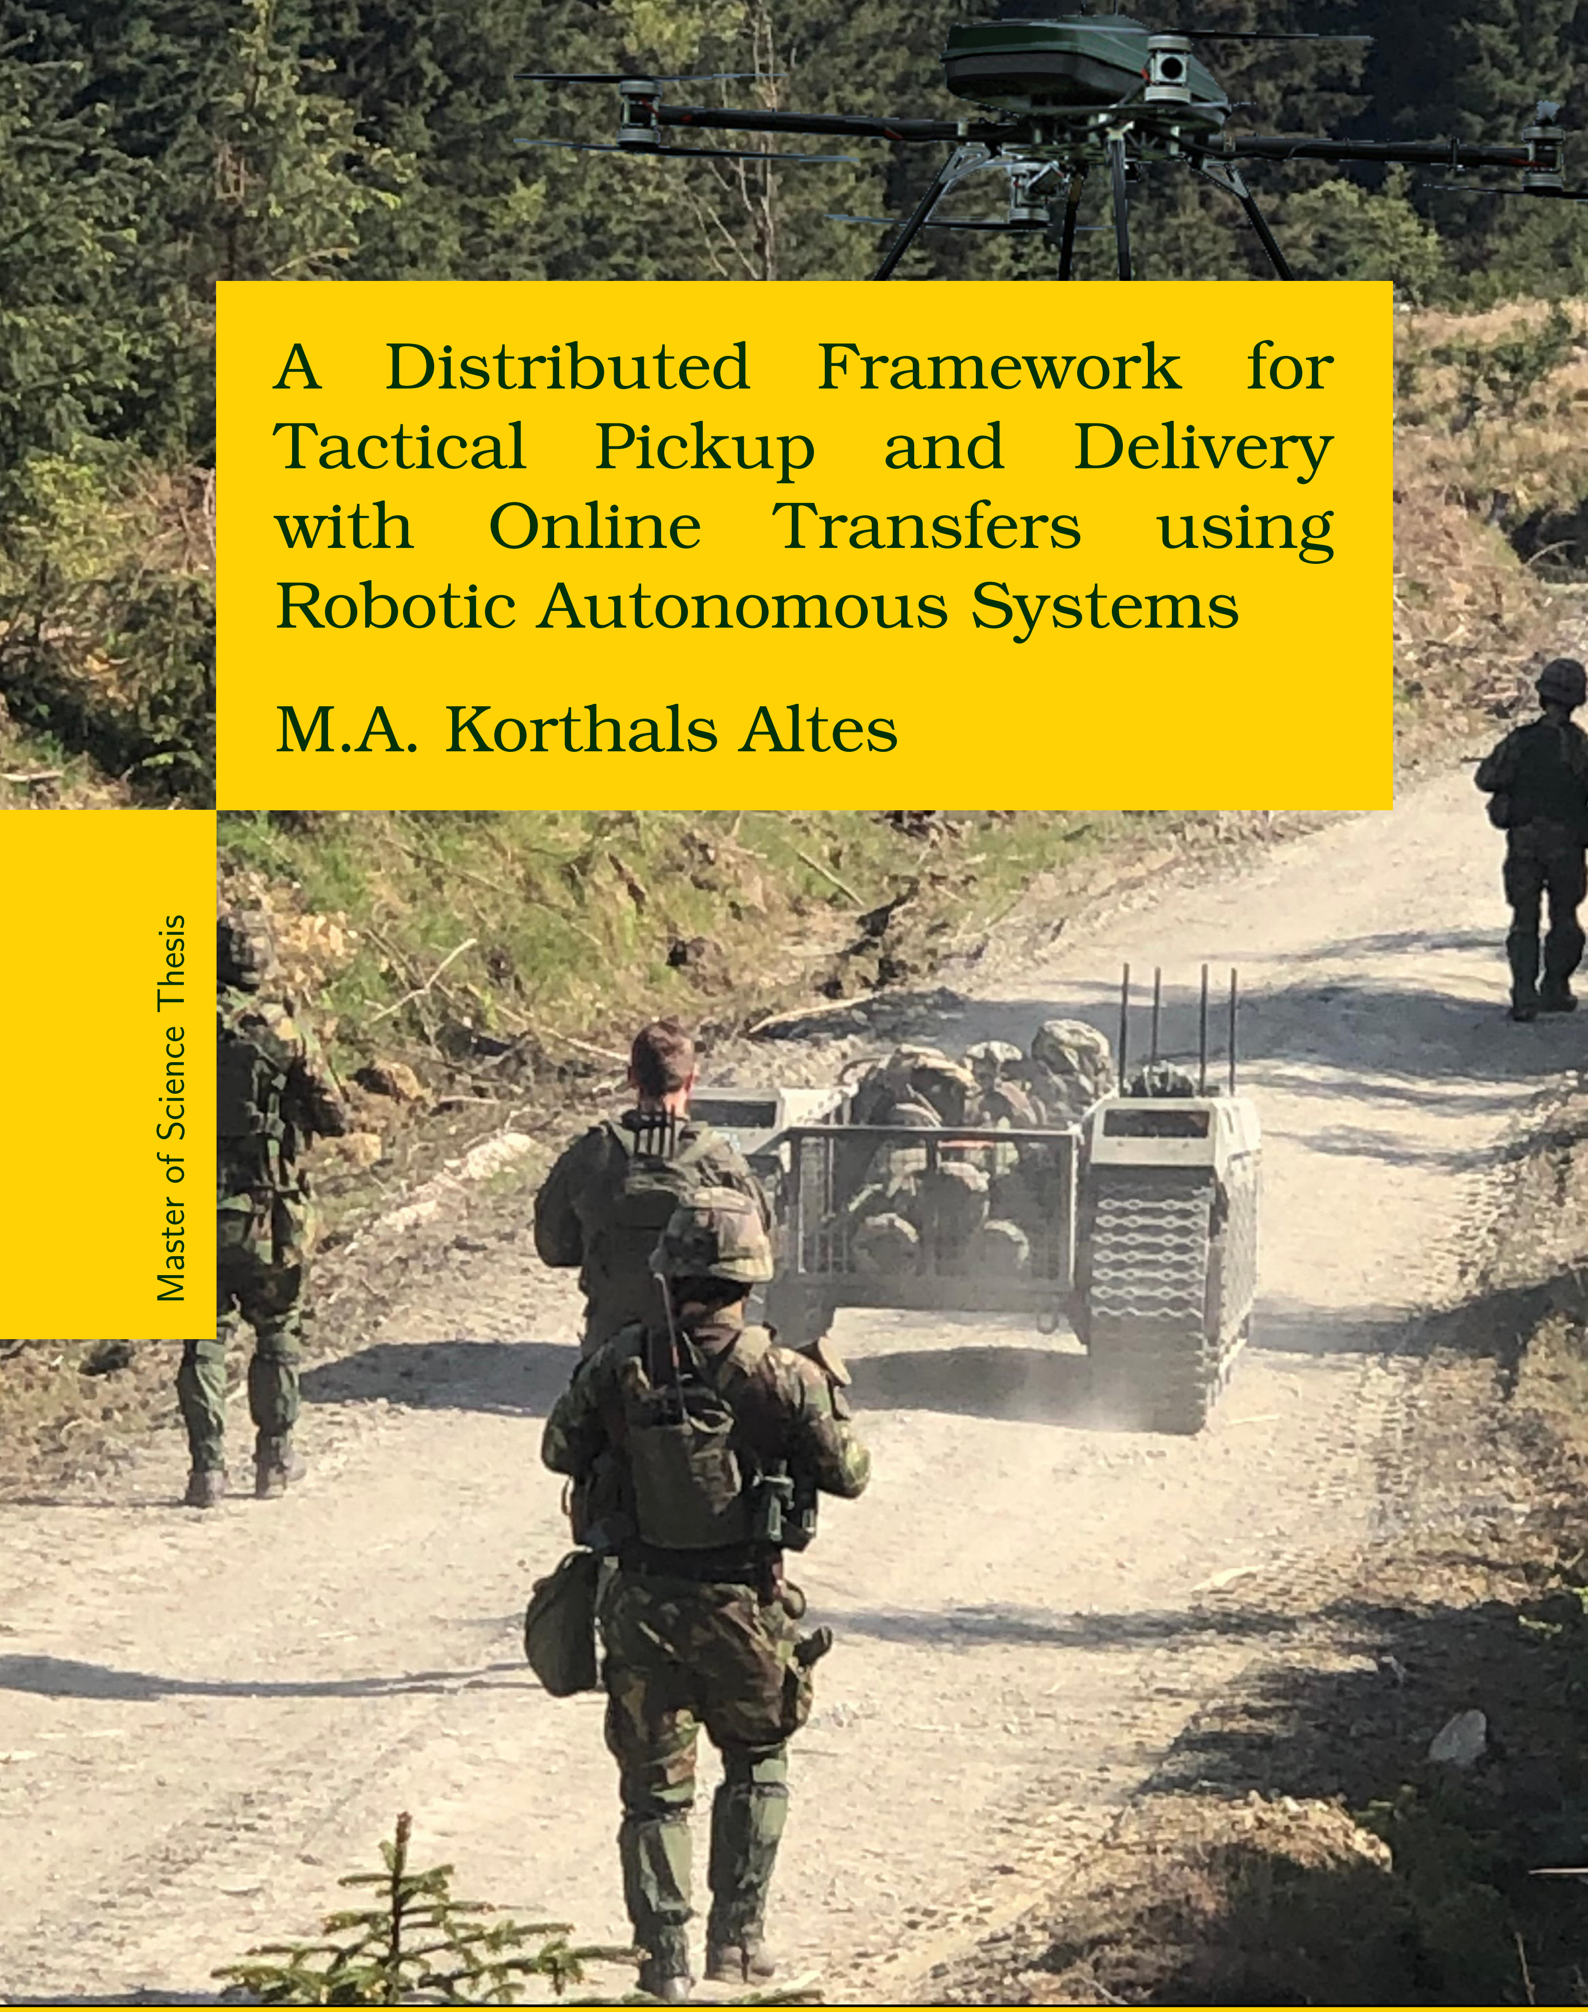

A Distributed Framework for Tactical Pickup and Delivery with Online Transfers using Robotic Autonomous Systems

M.A. Korthals Altes

Master of Science Thesis

Erasmus Universiteit Rotterdam

Erasmus School of Economics

A Distributed Framework for Tactical Pickup and Delivery with Online Transfers using Robotic Autonomous Systems

Master of Science Thesis

For the degree of Master of Science in Operations Research and Quantitative Logistics at Erasmus University Rotterdam

M.A. Korthals Altes

Student ID: 447875

Supervisor: dr. R. Spliet

Second Assessor: prof. dr A.P.M Wagelmans

November 23, 2020

Erasmus School of Economics (ESE) · Erasmus University Rotterdam

Koninklijke Landmacht

The work for this project was supported by the Royal Dutch Army. Their cooperation is hereby gratefully acknowledged.

ERASMUS UNIVERSITEIT ROTTERDAM

ERASMUS SCHOOL OF ECONOMICS

Abstract

The interest in using Robotic Autonomous Systems (RAS) capabilities has increased rapidly in the previous years, often with the idea of supporting new concepts of operations. One such concept is dispersed operations, for which the logistics pose significant challenges, making classical tactical resupply hardly a viable approach. Instead, a new RAS-based approach seems promising. As current research seems limited to autonomous Last-Mile Delivery (LMD) or static optimization models, an integrated approach named Dispersed Autonomous Resupply (DARE) is proposed.

The difficulty of creating a solution method for DARE is that it has to deal with a chaotic combat environment including risks of being destroyed, highly dynamic plans, and partial, unreliable information. To deal with these challenges, the Multi-Robot Systems (MRSs) based Distributed Constraint Optimization Problem (DCOP) paradigm is used. This allows for generating solutions locally on the involved RAS, while also enabling cooperation and coordination, such that scalability and robustness against unreliable communication can be achieved.

In this thesis, a distributed solution method is proposed. A two-stage heuristic is developed for solving the problem on each vehicle locally, which is combined with the DCOP-based algorithms MGM and MGM-2 for cooperation. By allowing cooperation, vehicles can synchronize their decisions and perform online supply transfers and recharging.

It is shown that a RAS-based setup for DARE is significantly more robust against the risk of destruction and quickly changing plans, showing more flexibility and survivability. In general, a RAS network not only appears to more consistently deliver supplies before they are needed, but it is also able to maintain the supply throughput better over a longer time.

Table of Contents

List of Figures

List of Tables

Preface

In the previous years, the Royal Dutch Army has set out an ambitious course of developing and integrating Robotic Autonomous Systems (RAS) capabilities in its operational concepts. I have been fortunate enough to be part of this development, alongside my knowledgeable and eager colleagues of the Defensity College program. Through this program I have also gained the complete freedom of developing novel RAS-based concepts, of which this thesis is the result.

Of course, I could not have completed this research on my own. I would like to thank my supervisor dr. R. (Remy) Spliet for his guidance and the continuous pushing to be more precise. Furthermore, I would like to thank Defensity College for enabling me to work on this project. Lastly, I want to thank my colleagues from the army, who gave purpose to this work.

Rotterdam, Erasmus University
November 23, 2020

M.A. Korthals Altes

"The essence of flexibility is in the mind of the commander; the substance of flexibility is in logistics."

— *Rear Admiral Henry Eccles, U.S. Navy*

Chapter 1

Introduction

1-1 Motivation

Due to a changing balance of global power and increased tensions at the border of the EU, the focus of NATO members has shifted from *bringing peace* to *keeping peace*. Inevitably, this renewed focus on the theme of high-intensity combat against a peer opponent in a large-scale conflict, of which a recent example is the ongoing confrontation between Ukraine and Russia. This conflict proves to be a critical evaluation of NATO concepts of operations, and especially that of the light brigade (13^{e} Lichte Brigade, 2018).

The new concept of operation prescribes nonlinear and dispersed combat, potentially supported by Robotic Autonomous Systems (RAS). Effectively this means that spatially distributed Combat Elements (CEs) conduct offensive operations in multiple directions at the tactical-operational level (Edwards, 2005). The essence of dispersed operations is that through enhanced Situational Awareness (SA) and computer-aided decision making the initiative can be placed at lower levels in the command hierarchy, enabling small CEs that operate independently while striving for *simultaneity* and *unity of effort* (Edwards, 2005), such that the state *converges* to the commander's intended end state (Alberts, 2007).

Dispersed operations, however, pose significant logistical challenges. As there is no classical, closed and connected front, there is no clear definition of a friendly or hostile area. Instead, distances between CEs increase, and uncovered areas increase in size and number which contributes to decreased and uncertain SA. Consequently, the risk of losing lives or supplies is increased as the supply chain possibly traverses hostile territory.

Naturally, it seems tempting to develop a RAS application to solve these challenges, but implementation is still lacking. One of the reasons is throughout literature the development of optimization models for tactical logistics seems to be treated as a separate technology compared to robotic platforms for Last-Mile Delivery (LMD), in which they are regarded as simple load carriers. Instead, it appears more promising to apply an integrated approach, in which a dedicated RAS-based concept of operation for tactical logistics is optimized.

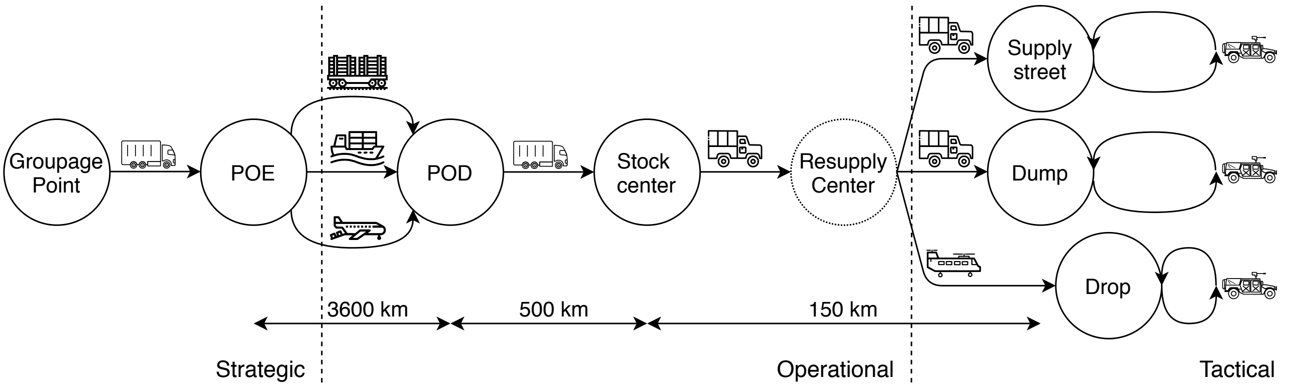

Figure 1-1: A schematic view depicting the flow of supply, based on Kablau (2002). Through different levels of combat, the supplies are moved towards the CEs. On the tactical level, an LCE is assigned to either place dumps, drop by air, or service through a supply-street.

1-2 Tactical logistics

Following the doctrinal definition of logistics by the US Marine Corps, "Logistics encompasses all actions required to move and maintain forces. This includes the acquisition and positioning of resources as well as the delivery of those resources. (MCDP-4)" (US Marine Corps, 2018). Logistics is split into three levels of war: strategic, operational, and tactical. As military tactics encompass the deployment and planning of forces on the actual battlefield, tactical logistics is concerned with employing Logistics Combat Elements (LCEs) designated to sustain the CEs. This concept is closely related to Combat Service Support (CSS), which describes the actual activity of providing services and supplies.

The Royal Dutch Army distinguishes five classes of supply, following the NATO standard. The LCE is tasked with supplying water and food (class I), munitions (class V), fuel (class III), and sometimes equipment (class II) or construction materials and spare parts (class IV). The classical flow of supplies is depicted schematically in Figure 1-1. This flow is based on the 'physical distribution' concept, as applied by the Royal Dutch Army (Kablau, 2002). The supplies flow from a strategic groupage point and Point of Embarkation (POE), entering the operational level at the Point of Debarkation (POD), after which it is shipped to one or more stock centers - the Main Operating Bases (MOBs). Possibly, also some forward resupply centers - also known as Forward Operating Bases (FOBs) - are set up, which are much smaller and more mobile. For the classical approach of tactical logistics, multiple methods exist to deliver the supplies to the CEs. The first is a supply-street (or so-called CSS area), in which CEs can pass through a mobile street of LCE vehicles, resupplying a certain supply class at each vehicle. As these supply streets are often designated as primary targets by enemy forces, they are located some distance behind the front and need to be sufficiently protected to ensure their security. Alternatively, the supplies can be left at a (hidden) dump, such that it can be picked up by CEs at a later point in time. Lastly, supplies can be brought closer towards the CEs by airdrop, although they will not be dropped onto fighting CEs themselves, such that units still need to retrieve the supplies.

The question remains if this classical approach for tactical logistics would suffice in a modern combat environment. Following the evaluation of NATO concepts of operation in the confrontation between Ukraine and Russia, the analysis of Karber (2015) states that an extremely efficient Reconnaissance Fire System (RFS) can be observed, by combining multiple, heterogeneous and layered Intelligence, Surveillance & Reconnaissance (ISR) capabilities with massive, concentrated and destructive firepower. Consequentially, the Royal Dutch Army identified that there is at least a range of 90 kilometers of unsafe terrain in which units can spend no longer than 2 to 3 hours without remaining unnoticed, such that units need to be constantly on the move (Koninklijke Landmacht, 2017). This conclusion effectively excludes the classical methods for tactical logistics as a viable approach. According to the Royal Dutch Army, the classical approach could instead be adapted to support dispersed operations in this new environment. It is suggested that this can be achieved by incorporating the LCEs within the maneuver elements (CEs) such that operational tempo can be increased (Koninklijke Landmacht, 2017). However, this would increase risk, drawing up increasing numbers of units to provide security. Though it is suggested that the increased use of sophisticated sensors for e.g. fire detection combined with smart algorithms that can automatically suggest maneuvers should negate some of these risks, it is admitted that a lot of extra armor and armament is still required. This would require enormous and unrealistic investments, as currently logistical systems are falling short in terrain accessibility, protection, armament, and Command & Control (C2) capabilities. As some early findings indicate, like the suggestion of Unmanned Areal Vehicle (UAV) predeployment to waiting areas (Koninklijke Landmacht, 2017), the possibility emerges for a completely novel tactical supply model based on RAS.

1-3 Dispersed Autonomous Resupply

Though its perceived benefits are clear, the question remains: How should the concept for tactical logistics be improved for dispersed combat operations by employing an integrated RAS-based approach? First, the performance indicators will be laid out which are used to demonstrate if a novel approach performs better than the classical approach. Furthermore, these indicators serve as a quantitative measure for the development of a solution method. Secondly, the novel concept of operations will be formulated, which will be identified as Dispersed Autonomous Resupply (DARE). The involved RAS then compose a DARE network.

1-3-1 Performance indicators

Based on existing logistical doctrines, five important principles for tactical logistics can be deduced (US Marine Corps, 2016). The first three principles describe the effectiveness of the logistical process, which needs to be balanced with the fourth principle, efficiency. Survivability is needed to maintain both effectiveness and efficiency. Ineffectiveness could result in disastrous consequences for any operation. Therefore the goal is maximizing effectiveness in the most efficient manner. An effort is made to quantify the principles for tactical logistics, resulting in the corresponding performance indicators.

1. **Responsiveness:** The right support at the right place at the right time. This means that the CEs should not be hampered in their decision making by logistics.

   → *Indicator*: Providing the required supplies before an action is planned to be executed by any CE, effectively minimizing delay on deliveries.

2. **Sustainability:** The ability to maintain support during the whole operation for all CEs involved. This includes an effective throughput of supplies even during peak consumption.

   → *Indicator:* The supply level of each CE should be kept continuously at a sufficient level to maintain its operations.

3. **Flexibility:** The ability to adapt to changing circumstances and unforeseen events. This can be achieved by anticipation, redundancy, reserves and effective C2.

   → *Indicator:* Robustness against sudden change of plans, measured by the degree to which responsiveness and sustainability can be maintained after plan alterations.

4. **Survivability:** The ability to maintain the required responsiveness and sustainment even when vehicles are destroyed.

   → *Indicator:* Robustness against destruction, meaning that under varying threats the effectiveness should still remain adequate.

5. **Efficiency:** The ability to provide sufficient support at the least costs, both from an economical and ethical perspective, which means that unnecessary risk should not be taken and that an excess in resources is undesirable due the financial burden and general scarceness of supplies.

   → *Indicator:* As the fuel consumption is assumed negligible compared to the cost of losses, efficiency is measured in the total loss of vehicles and supplies during resupply.

1-3-2 Concept of operation

Before a RAS-based concept of operation can be laid out, some basic assumptions must be made about the vehicles involved such that the principles can be satisfied:

- All involved vehicles have the ability to navigate through and traverse difficult terrain types, where different vehicles have different terrain accessibility.
- The vehicles have the ability to hide, or at least choose a path maximizing cover.
- The low-level controllers are included on the vehicle.
- The ability for inter-system cooperation physically (transfer of supplies) and cognitively (joint optimization). This is achieved by active communication links (albeit not necessarily trustworthy), local computation power and robotic arms for transferring supplies.

Environment A DARE network would be deployed when tactical logistics needs to be provided for the dispersed operations of multiple platoon-sized CEs. Furthermore, it is known that a peer-opponent can deploy an effective RFS (see Section 1-2). At any time, the systems can therefore be hit by a strike, for which the probability is increased if systems remain in a fixed position longer. Furthermore, the systems needs to resupply over long distances, as the closest FOB is located at least 90km from opposing forces. Lastly, specific terrain types are not accessible by all systems.

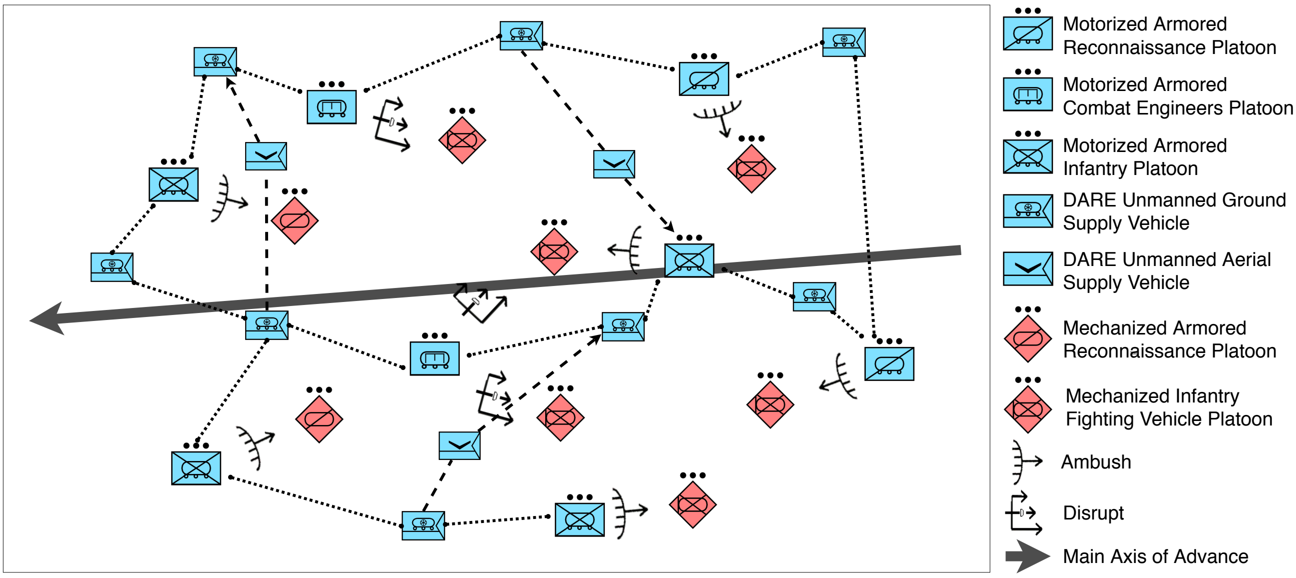

Figure 1-2: A figurative operation sketch for DARE. Dotted lines represent the closest neighboring UGV or CEs. The dashed lines represent the routes of dispatched UAVs transporting supplies from UGVs to CEs or other UGVs. It becomes clear the UGVs choose their position to optimally sustain the combat operations while minimizing their own risk to be eliminated.

Network The DARE network consists of multiple heterogeneous systems, including both small and large Unmanned Ground Vehicles (UGVs) and UAVs. The systems are supplied from the nearest FOB, and supply the CEs. All vehicles have the capacity to store supplies, such that they can act as a capacitated, mobile depot.

Objective The objective of the DARE network is to make decisions such that the resupply effectiveness can be maximized at the least costs (see Section 1-3-1). This effectiveness is maximized by enlarging the expected value for each of the effectiveness indicators, while supply losses are minimized. This can be achieved by increasing redundancy and decreasing the risk taken by each vehicle.

Process It is required that plans or (computed-aided) predictions for maneuvers, supply consumption and (unexpected) combat are provided digitally to the DARE network as soon as received by logistical officers. Effectively, large amounts of supplies are then moved based on these high-level directives. Then, the network can make continuous decisions to optimally support the predicted and planned operations with maximum effectiveness. The network can achieve this by continuous repositioning and reallocation of supplies. In Figure 1-2 a figurative operation sketch of the process is provided. The dispersed CEs either disrupt or ambush the hostile forces while they move over their main axis of advance. The UGVs reposition and reallocate supplies such that the actions of the CEs are anticipated. The UAVs accompanying the UGVs are used as a quick dispatcher to either deliver supplies to CEs, or reallocate them to UGVs.

1-4 Problem statement

The difficulty with generating solutions for a DARE network is that the operating environment is plagued by high uncertainty, many communication constraints and lacking information, making the development of a solution method a complex challenge. Furthermore, a concept that relies on both flexibility and redundancy requires scalability, which poses problems for the tractability of a solution method. This thesis is focused on exploring the possibility of modeling DARE based on existing modeling frameworks used in Operations Research and developing specific distributed algorithms needed to obtain a quality solution. The corresponding challenges for the modeling and solution approach can be summarized as:

- A model that describes the DARE concept, such that effectiveness can be maximized efficiently.
- Dealing with a stochastic environment in which unexpected events may occur that influence demand for specific types of supply at unexpected times and locations.
- A highly dynamic environment, in which decisions might need to be made based on new information within minutes. The solution approach should inhabit the flexibility to quickly respond to this new information.
- Locally only partial information through unreliable communication links is available.
- Heterogeneous systems varying in capacity, speed, battery capacity, protection, signature and terrain accessibility.

These points indicate that a model should be developed with a corresponding solution algorithm that is both fast, scales well, provides good solutions, and deals with information uncertainty efficiently. The central research question can therefore be described as:

Can an efficient solution for DARE be developed such that dispersed combat operations can be effectively supported while remaining scalable in computation time and sufficiently robust against unreliable communication?

1-5 Overview

In Chapter 2 the literature aimed at improving military tactical logistics will be discussed. Furthermore, related modeling methods and the corresponding solution methods are discussed. In Chapter 3 a problem description is given by formalizing a rigorous mathematical model describing DARE. In Chapter 4 the methodology is laid out. First, the solution structure is described including coordination, cooperation and the local problem. Then, the chosen local and distributed optimization algorithms will be discussed. In Chapter 5 a scenario is provided for data generation, after which in Chapter 6 the solution methods are analyzed based on subsequent simulations. In Chapter 7 overarching conclusions will be drawn, along with a critical discussion of the proposed methods.

Chapter 2

Literature Review

2-1 Military oriented improvements

Ivanova, Gallasch, and Jordans (2016) provide an elaborate report on the use of autonomous systems for Combat Service Support (CSS). Their main contribution is listing important innovations, for which a technology prioritization is formulated in cooperation with the Australian Department of Defence (DoD). From this, two distinct categories emerge:

- Smart analytics: the ability to include prediction and optimization in different decision tools to increase sustainment at lower costs.
- Autonomous platforms: The ability to perform Last-Mile Delivery (LMD) through complex terrain with Unmanned Areal Vehicles (UAVs) or Unmanned Ground Vehicles (UGVs) or employ long-range convoys by unmanned platooning, to reduce both costs and risk for personnel.

Notable is that they treat analytical tools as a separate technology from robotic platforms, apart from proposing strategic predeployment of UGVs (though no source is included). As this view is found to be surprisingly common throughout literature, it will be discussed following this prioritization. First, optimization models improving the logistical process are discussed. Second, examples of improved (partially) autonomous platforms are provided.

2-1-1 Optimization models

A basic approach is taking existing models for scheduling or inventory optimization but include extra military constraints. Sebbah, Ghanmi, and Boukhtouta (2011) consider a Vehicle Routing Problem (VRP)-like model including heterogeneous supplies and heterogeneous vehicles, with additional military constraints. Sebbah, Ghanmi, and Boukhtouta (2013) extend this further by also optimizing for the optimal fleet mix and size. Alternatively, Ren, Zheng, and Tan (2013) optimize inventory based on the multi-item newsboy problem, by including both stochastic demand, and stochastic supply due to enemy attacks. Furthermore, Gallasch et al. (2008) propose to model a military logistics network using a Coloured Petri Net (CPN), such that delivery is triggered through the network after passing certain inventory thresholds.

These military approaches can be combined to optimize both inventory and scheduling jointly. Baker and Shi (2002) consider a network topology with a continuous flow of supplies. Assuming quadratic cost relationships for protecting supplies in depots or during transport, an optimal control policy can be derived, aiming to minimize cost while satisfying demand. A discrete approach is suggested by Marufuzzaman, Nur, Bednar, and Cowan (2020), who consider an extended Multi-Commodity Network Design Problem (MCNDP). They include demand satisfaction subject to a certain priority and include dynamic risks for utilizing certain routes. Though both approaches include some kind of cost for protection, no mobility of Combat Elements (CEs) or depots is allowed.

Mobility of supplies is considered by Barahona et al. (2007), who focus on the distribution of spare-parts in Network-Centric Operations (NCOs). As input, a forecast on breakages is used in combination with information on the current inventory levels at CEs and their planned movements. A two-stage solution approach first sets the desired inventory levels and schedules the redistribution afterward. Though mobility plays a key role, it is not an active decision. Instead, Toyoglu, Karasan, and Kara (2011) propose a mobile Ammunition Distribution System (ADS), comparable to the Location Routing Problem (LRP). The basis is a combination of fixed transfer points and mobile transfer points, with the objective to deliver ammunition as close to units as possible, within a specific time window. Based on risks different objectives like lower driving time versus lower costs could be achieved. Gue (2003) proposes the most extensive model, named the dynamic distribution problem, which appears comparable to the Inventory Routing Problem (IRP). The objective is to minimize shore-based inventory when performing Marine Corps operations supported from the sea. A ship is considered a resupply center, from which onshore depots can be supplied. CE movements are preplanned and known in advance, as are the possible locations for mainland depots. From the depots, different vehicles can be scheduled to supply CEs. Furthermore, depots can move to better support operations. Lastly, these depots can have varying inventory, which is to be minimized while still satisfying demand over the complete time horizon. For both approaches, however, transfers remain hierarchical and the risk of occupying certain locations is not taken into account directly.

2-1-2 Autonomous systems

The second category consists of literature and projects dedicated to improving LMD or automated convoys. LMD describes the last link in the tactical supply chain from either a dump, drop or supply street to the CEs, as is displayed in Figure 1-1. As noted, LMD does not consist of a mile, but can even span up to 90 kilometers (Koninklijke Landmacht, 2017). Though being a hot topic in military innovation, it remains limited to separate development of hardware systems and optimization models. For example, the UK Ministry of Defence (2017) propose scenario vignettes in which a UGV supplies spare parts and fuel over a distance of 30km to a broken Infantry Fighting Vehicle (IFV), and UAVs supplying an infantry platoon with emergency ammunition over a distance of 2km in heavy weather. Clearly, only the use of autonomous systems as quick, easy, and cheap dispatcher is discussed, without the flexibility for optimization.

Other initiatives show the same perspective. Ackerman (2014) writes about the promising results of autonomous driving during trials with unmanned truck convoys. Furthermore, he writes about initiatives with disposable UAVs (Ackerman, 2017) and precision supply drops (Ackerman, 2019). Other promising results are achieved by Qinetiq as part of the UK competition with a Milrem Themis capable of driving off-road autonomously (Walker, 2018). Furthermore, Keirin (2016) considered the effects of unmanned systems and the consequential adaptations to modern logistical software. Though seen through a perspective of operations research, it does not provide an alternative modeling framework for autonomous logistics, but merely considers the effects of certain hardware systems. Thornton and Gallasch (2018) analyzed the use of swarms for LMD, in which they analyze the benefits and vulnerabilities of using swarms as opposed to single vehicles. However, though multiple vehicles are considered, delivery is still performed as a one-off mission, instead of multiple vehicles cooperating autonomously for a longer duration of time.

2-2 Related modeling methods

Apart from military-oriented improvements, there is also a wide range of models that are related based on similar objectives and constraints. The purpose of this section is to provide a broader view of the complexity and possibilities of those models. A model is considered related if it either describes some form of storage at mobile depots, pickup, delivery and transfer of supplies, the balancing of multiple types of cost, or those who consider relocation or demand in a dynamic, stochastic time-extended environment.

An example of relocation for varying demand within a time horizon is provided by the Dynamic Facility Location Problem (DFLP) (Farahani, Abedian, & Sharahi, 2009), which is closely related to the Maximum Coverage Facility Location Problem (MCFLP) (Church & ReVelle, 1974). This model can be combined with a VRP, resulting in the LRP (Klibi, Lasalle, Martel, & Ichoua, 2010). In general, there is a maximum amount of vehicles, which are shared over all possible depots (Prodhon & Prins, 2014). This problem is adopted to a military context for ammunition distribution by Toyoglu et al. (2011). In the context of LMD, Chauhan, Unnikrishnan, and Figliozzi (2019) attempted to combine the problem with drones.

Though relocation is discussed, transfers are not allowed. Models in which transfers occur frequently are those that combine delivery with vehicles and UAVs. An example is the Flying Sidekick Traveling Salesman Problem (FSTSP), in which vehicles can be interpreted as mobile depots (Murray & Chu, 2015) (Dell'Amico, Montemanni, & Novellani, 2019). Possibly, the launch locations are fixed (Ferrandez, Harbison, Weber, Sturges, & Rich, 2016). A comparable problem is provided by the Vehicle Routing Problem with Drones (VRPD) (Wang, Poikonen, & Golden, 2017) or the Traveling Salesman Problem with Drone (TSPD) (Agatz, Bouman, & Schmidt, 2018). Many of these solution approaches consider the limited energy capacity of drones (Dorling, Heinrichs, Messier, & Magierowski, 2016) (Shavarani, Nejad, Rismanchian, & Izbirak, 2018) (Rabta, Wankmüller, & Reiner, 2018) (Jeong, Song, & Lee, 2019). Carrying a sidekick is generalized further in the Vehicle Routing Problem with Trailers and Transshipment (VRPTT), focusing on synchronization constraints (Drexl, 2013).

Pickup and delivery for either paired or unpaired goods are combined in the Vehicle Routing Problem with Pickup and Delivery (VRPPD) (Parragh, Doerner, & Hartl, 2008). Supplies are classified as unpaired, which means that after pickup it can be used to satisfy any demand. The Parallel Drone Scheduling Traveling Salesman problem (PDSTSP) combines this with UAVs (Ham, 2018). By also allowing transfers the more flexible Pickup-and-Delivery Problem with Transfers (PDPT) is created. This larger degree of freedom should result in greater efficiency (Masson, Lehuédé, & Péton, 2013), especially in high demand situations (Berbeglia, Cordeau, & Laporte, 2010). Bouros, Sacharidis, Dalamagas, and Sellis (2011) extend this by considering the dynamic variant in which requests arrive at arbitrary times. However, transfers are assumed to happen by leaving the goods at a transfer point. To allow transfers between vehicles, the Pickup-and-Delivery Problem with Online Transfers (PDPOT) is introduced. To synchronize these transfers, dynamic time windows are needed (Drexl, 2012). For offline transfers, synchronization requires precedence (a delivery at a transfer point precedes the pickup). In the case of online transfers, the synchronization is exact, and pickup and delivery occur simultaneously. Though this cooperative approach can result in benefits, possible bottlenecks can occur when trying to enforce synchronization (Otto, Agatz, Campbell, Golden, & Pesch, 2018). Additionally, the transfer point is not even discrete, as with for example mid-air refueling of aircraft (Coltin & Veloso, 2014).

Prior models only maximize some form of demand metric. In contrast, demand and inventory levels are considered jointly in the IRP (Bertazzi, Savelsbergh, & Speranza, 2008). Especially interesting is the multi-depot case with stochastic demand (Roldán, Basagoiti, & Coelho, 2017). This problem is classified for tactical logistics as the Dynamic Distribution Problem by Gue (2003). Instead of inventory, the MCNDP jointly optimizes the selection of arcs to operate and the corresponding commodity flow (Fragkos, Cordeau, & Jans, 2017). Marufuzzaman et al. (2020) use this to describe risk over certain arcs.

All of the discussed modeling approaches described in some way a part of the Dispersed Autonomous Resupply (DARE) problem. Summarizing, DARE could be classified as a combined instance of MCNDP for taking into account risk over different arcs, with the PDPOT allowing for instantaneous transfers between systems, the VRPTT or FSTSP for the combination UAVs and UGVs, and a dynamic LRP for deciding at which locations supplies are held.

2-3 Solution approaches

Evidently, the problem at hand is of extreme complexity, meaning that generating solutions will be a difficult matter. Furthermore, the generation of solutions itself is subject to a challenging military environment, complicating matters even further. The question remains what general solution approaches for such a complex problem are viable for this environment. Therefore the different centralized and distributed approaches are discussed in relation to the requirements listed in Section 1-4.

2-3-1 Centralized

In all of the mentioned papers in Section 2-1 and 2-2, the proposed formulations are solved centrally. Though many different approaches are discussed, solutions are found using either exact algorithms, meta-heuristics, heuristics, or a combination. Apart from computational tractability, the major issue with all of the above approaches is that it is not applicable to the military environment. As described in Section 1-4, locally only partial information through unreliable links is available. Furthermore, there is a need for quick decision-making in a highly dynamic environment. This eliminates the centralized approach, as this requires all information to be trustworthy. In an unreliable network, such a task is cumbersome if not impossible. Given this was successful, the computation time is often still too long, and even if it were on time, it cannot be trusted that the solution would reach the actual system that needs to perform a task.

2-3-2 Distributed cooperation

Opposite to centralized is the distributed approach. Two types of distribution are distinguished. Firstly, distribution can exploit parallelism (like the ones arising from decomposition), to enhance performance and scalability. In practice, this is often decentralization, in which multiple non-cooperating computation sources are active in parallel and a central source performs synchronization. The second class arises from Multi-Agent Systems (MASs), or more specifically Multi-Robot Systems (MRSs). Here, the distribution is natural as information is present locally and no central agent is available to compute the global optimum. For truly MAS-based distributed solution approaches, some form of cooperation is necessary to improve the solution quality. Cooperation can be categorized into three types:

- Passive: Each agent optimizes its problem individually, but agents communicate their local context, such that agents show collective behavior by reacting to each other.
- Implicit: Each agent optimizes its problem individually, based on planned decisions of other agents, as this is communicated with the local context.
- Explicit: Each agent cooperatively optimizes the problem to obtain a globally optimal solution, by passing messages depending on the algorithm choice.

An important framework for implementing explicit cooperation is a Distributed Constraint Optimization Problem (DCOP), which describes a problem in which agents must globally optimize a set of variables, while subjected to distinct sets of constraints. In this case, agents are responsible for their own variables which represent local (physical) behavior.

For all types of cooperation, however, communication is necessary to maintain performance. Therefore the available solution methods do not fulfill the requirements from Section 1-4 completely, as there is no guaranteed robustness against communication failures.

2-3-3 Distributed coordination

To overcome the dependency on communication arising from the solely cooperative approach, a coordinative approach could return better results for MRSs. For instance, Skubch (2013) analyzes nearly identical requirements as in Section 1-4 and proposes coordinated teamwork. The idea is that when agents commit themselves in an efficient manner of working to a specific goal, robustness is increased when communication is failing. Another influential work even originates from a military problem, and resulted in the concept of joint intentions and shared plans (Tambe, 1997). An important hypothesis by Tambe (1997) is that completely preplanned coordination would not be flexible enough, leading to drastic failures (much resembling the centralized approach under dynamic circumstances). Therefore some dynamic form of coordination needs to be applied to be able to flexibly adjust to new information under realistic communication constraints. This type of coordination can take place by implementing a kind of organizational structure in the MRS. A review of these different structures is provided by Abbas, Shaheen, and Amin (2015). They note that “the organizational level describes the “what” and not the “how”. In other words, the organizational level imposes a structure into the pattern of agents’ activities”. Effectively, coordination imposes limits on the solutions generated by the agents.

Based on this concept of coordination, a general framework is proposed based on contemporary military intent-based Command & Control (C2) (Korthals Altes, 2021). Intent-based C2 is based on the idea that at any level in a command hierarchy a commander should have the freedom to exploit local advantages without explicit approval from superiors, given the superiors’ intent is completely clear. The assumption is that local initiatives will then also contribute to the global objectives, even if communication is failing or delayed, as it can be assessed locally if the decisions are in line with the global objectives. Concluding, cooperation remains vital for performance, but in the case of cooperation being not possible due to failing communication, robustness is guaranteed through effective coordination.

Chapter 3

Modelling DARE

In this chapter, a description is given of the model used to represent the Dispersed Autonomous Resupply (DARE) network.

The environment is represented by a directed graph $G(N, A)$ with nodes N and arcs A. The nodes describe physical locations, and the arcs describe all paths between these locations. Furthermore, the model is extended over a horizon of T time periods.

Within this environment, a set $\mathcal{C}$ of Combat Elements (CEs) conduct their operations. Their planned positions are provided for all $c \in \mathcal{C},\ i \in N,\ t \in T$, such that $y_{cit} = 1$ if the CE is present, and $y_{cit} = 0$ otherwise. Over time, these CEs consume different types of supplies, which are grouped in classes defined by the set K. Their estimated supply consumption in kilograms of each class is given by C_{kct}, and the corresponding inventory level by I_{kct} $\forall k \in K$. Furthermore, also a set $\mathcal{F}$ of Forward Operating Bases (FOBs) is provided, for which the fixed positions are described in an equal manner by y_{fi} $\forall f \in \mathcal{F}$. These FOBs act as depots, at which vehicles can be recharged and supplied indefinitely. After each time period passes, new plans for the CEs can be provided to the model, which do not necessarily have to correspond to the previously provided plans.

To bridge the distance from the FOBs to the CEs, different Robotic Autonomous Systems (RAS) are available to pick up and deliver supplies. A set $\mathcal{D}^g$ of Unmanned Ground Vehicles (UGVs), and a set $\mathcal{D}^a$ of Unmanned Areal Vehicles (UAVs) compose the set of vehicles $\mathcal{D} = \mathcal{D}^g \cup \mathcal{D}^a$, each carrying an inventory I_{kdt} $\forall k \in K$. A vehicle can decide to either be present at a location $i \in N$ at time $t \in T$, meaning $y_{dit} = 1$, or decide to relocate from one location to another, defined by the decision $x_{dijt} = 1$ $\forall \{i, j\} \in N$. This relocation takes t_{dij} time, after which it can be present again at the arrival node, or decide to relocate immediately. This is clarified in Figure 3-1. The UGV decides to relocate at the start of period 1 and relocate in period 2 immediately. When arriving in period 3, it remains present such that it can interact with the CE that is present in period 3 too. However, in period 4 it may not interact with the UAV, as it has decided to relocate again.

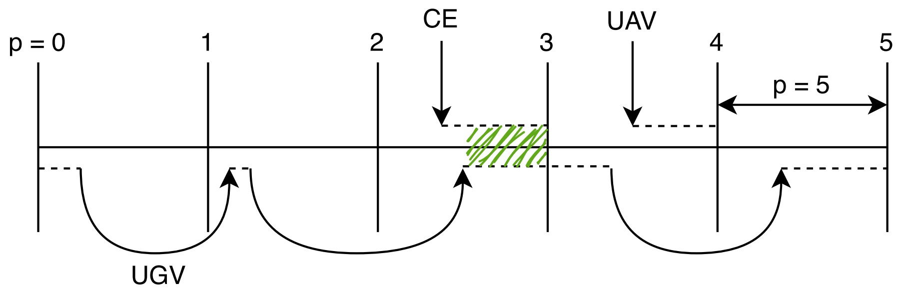

Figure 3-1: A depiction of a UGV traveling over a span of 5 time periods. It may interact with the CE in period 3, but not with the UAV in period 4, as it has decided to relocate instead of remaining present.

To be able to travel, each vehicle $d \in \mathcal{D}$ carries an energy level E_{dt}, limited by the energy capacity r_d. During movement, vehicles consume energy when traversing an arc given by $C_{dij}\ \forall i, j \in N$, such that the energy level is updated in each period:

$$E_{d,t+1} = E_{dt} + R_{dt} + \sum_{d' \in \mathcal{D}} R_{dd't} - \sum_{i,j \in N} C_{dij} x_{dijt} \qquad \forall d \in \mathcal{D},\ t \in T \tag{3-1}$$

This equation states that the energy level for vehicle $d \in \mathcal{D}$ is decreased by the decision to relocate to a new location, and increased by deciding to recharge an amount of energy $R_{dd't}$ from vehicle $d' \in \mathcal{D}$ or R_{dt} at any FOB. Conversely, this means that $R_{dd't} = -R_{d'dt}$.

Using the energy to move around, the vehicles can make multiple decisions. First, each vehicle $d \in \mathcal{D}$ can decide to pick up an amount of supply P_{kdt} at time $t \in T$ for each class $k \in K$ at any FOB $f \in \mathcal{F}$, as long as they are both present at that location at that time, meaning $y_{dit} = y_{fi} = 1$. The inventory for each class can be increased up to the vehicle's class capacity b_{kd}. Secondly, each vehicle can decide to deliver an amount of supply for each class $k \in K$ at time $t \in T$ to a CE $c \in \mathcal{C}$, denoted by D_{kdct}. Again, both must be present at the same place, meaning $y_{dit} = y_{cit} = 1$. Furthermore, delivery is limited by the available vehicle's inventory I_{kdt} and the capacity for each class of the CE b_{kc}. It is not required for a vehicle to satisfy all demand of a CE, such that the vehicle might hold reserves for other units or future consumption. Thirdly, vehicles $d, d' \in \mathcal{D}$ can decide to transfer an amount of supply of each class between them at a certain time, denoted by $T_{kdd't}$, but only if both vehicles are present at the same location, such that $y_{dit} = y_{d'it} = 1$. Transferring implies that $T_{kdd't} = -T_{kd'dt}$, as is the case with recharging. Combined, the evolution of the inventory of a vehicle is given as:

$$I_{kd,t+1} = I_{kdt} - \sum_{c \in \mathcal{C}} D_{kdct} + P_{kdt} + \sum_{d' \in \mathcal{D}} T_{kdd't} \qquad \forall k \in K,\ d \in \mathcal{D},\ t \in T \tag{3-2}$$

And the inventory of a CE as:

$$I_{kc,t+1} = I_{kct} - C_{kct} + \sum_{d \in \mathcal{D}} D_{kdct} \qquad \forall k \in K,\ c \in \mathcal{C},\ t \in T \tag{3-3}$$

All of the relocation, pickup, delivery, transfer, and coupling decisions are made to optimally supply the CEs. To model this, a threshold is defined for each supply class $k \in K$ at each time instance $t \in T$ for each CE $c \in \mathcal{C}$ as the minimum supplies necessary to operate, denoted

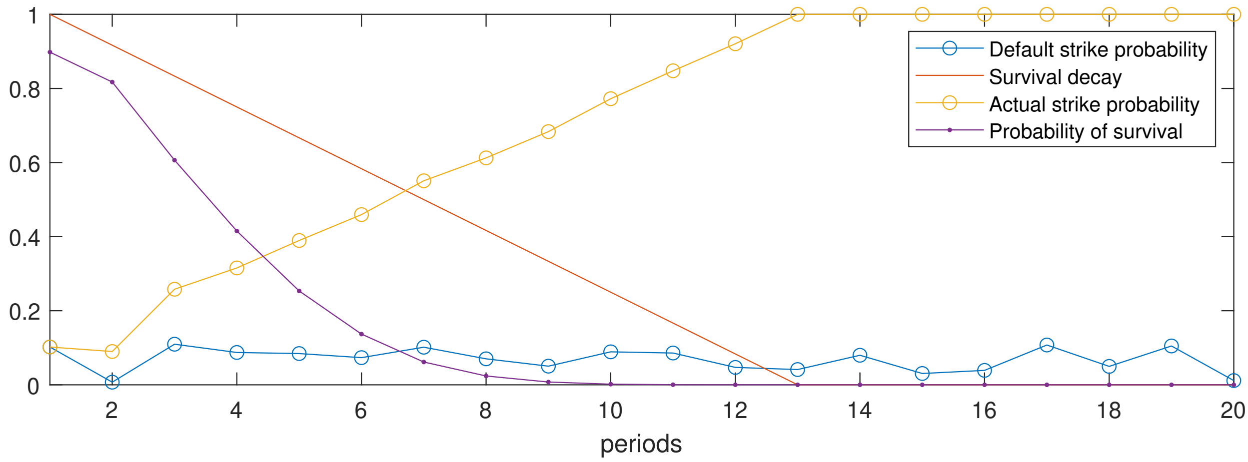

Figure 3-2: The actual strike probability increases due to a linear survival decay when remaining static for multiple consecutive periods. This results in a rapid decrease of the cumulative probability for survival.

by τ_{kct}. If this threshold is violated, a shortage is present, meaning $I_{kct} < \tau_{kct}$. This shortage is penalized according to provided weights w_{kct}, such that certain supply classes for certain CEs at specific times can be prioritized. Then, the objective is defined as minimizing the maximum shortage penalty, in order to lower the worst-case outcome over all CEs:

$$\min \max_{\forall k \in K, c \in \mathcal{C}, t \in T} (w_{kct} \cdot (\tau_{kct} - I_{kct})^+) \tag{3-4}$$

Furthermore, there is always a risk for enemy strikes during both relocation and static presence. If a strike occurs, the vehicle is destroyed and cannot perform deliveries anymore.

The default risk for an enemy strike at node $i \in N$ at time $t \in T$ when vehicle $d \in \mathcal{D}$ remains static is denoted as $\mathbb{P}^0_{dit}$. This risk could be based on for instance known enemy weapons, estimated locations, and the profile of the vehicles. The estimation is performed upfront, and provided to the model. If the vehicle remains present in the same location for n consecutive time periods, the static probability for a strike will gradually increase following the survival decay function $\phi(n_{dit})$. As described in Section 1-1, a unit can remain static no longer than 3 hours without being noticed. Hence, regardless of the function used, the risk for a strike is maximized after 3 hours, and chances for survival are approximately zero. Put otherwise, with t_p as period length, after $\bar{n} = \frac{3h}{t_p}$ periods the risk is 100%. Then, using an affine survival decay starting after the first period:

$$\phi(n_{dit}) = \max(1 - \frac{1}{\bar{n}}(n_{dit} - 1), 0) \qquad \forall d \in \mathcal{D},\ i \in N,\ t \in T \tag{3-5}$$

Then the probability for a strike depending on the periods a vehicle remains static equals:

$$\mathbb{P}_{dit} = 1 - (1 - \mathbb{P}^0_{dit}) \cdot \phi(n_{dit}) \qquad \forall d \in \mathcal{D},\ i \in N,\ t \in T \tag{3-6}$$

In Figure 3-2 this is displayed for a single system at a specific position. Given the vehicle has arrived in period 0 and each period constitutes 15 minutes, the survival chance has decayed completely after 12 consecutive periods. As is visible, the actual strike probability is not monotonically increasing as survival chances decay. This is due to the varying default probability for that specific position over time. For example in period 2, the default probability is that much lower than in period 1, that even though chances for survival decay, the actual strike probability is still lower. After the full 3 hours, however, a strike is imminent.

Furthermore, when a vehicle relocates to a new location it is exposed, such that the probability for a strike at vehicle $d \in \mathcal{D}$ while starting to move from node $i \in N$ to $j \in N$ at time $t \in T$ is described by $\mathbb{P}_{dijt}$. Combined, this means that a future probability for survival for each vehicle at time t is estimated based on the planned relocations:

$$\mathbb{P}^s_{dt} = \prod_0^t \left(1 - \sum_{i \in N} \mathbb{P}_{dit} \cdot y_{dit}\right)\left(1 - \sum_{i \neq j \in N} \mathbb{P}_{dijt} \cdot x_{dijt}\right) \qquad \forall d \in \mathcal{D} \tag{3-7}$$

Effectively, this means that future inventory levels depend on the probability of survival of the vehicles delivering supplies, such that the objective of a maximum shortage penalty is influenced negatively. Therefore avoiding risks increases the expected amount of supplies delivered, and sufficient redundancy would need to be included to compensate for each other's limited chances of survival.

Chapter 4

Methodology

4-1 Solution structure

In this section, the overall structure of the solution process is explained, highlighting the dependence between coordination, cooperation, and the local problem of each vehicle.

4-1-1 Coordination

To increase scalability and robustness against communication failures a distributed hierarchical intent-based Command & Control (C2) framework (Korthals Altes, 2021) can be implemented, which allows for effective coordination (see Section 2-3-3). This framework functions by dynamically partitioning the problem space into smaller subsets, which are assigned to individual agents, such that the problem can be solved locally.

In Figure 4-1 the effect of partitioning on the network topology is displayed. One can see the difference between the centralized approach, a distributed fully connected one, and a coordinated approach. The example shows the concept of overlapping subsets after coordination, such that each vehicle has a specific set of neighbors.

For the model described in Section 3, coordination can take place on the following elements:

- Limiting the traversable nodes to the subset $N_d \; \forall d \in \mathcal{D}$.
- Limiting the set of vehicles for cooperation $\mathcal{D}$ to $\mathcal{D}_d \; \forall d \in \mathcal{D}$
- Limiting the available Forward Operating Bases (FOBs) $\mathcal{F}$ to $\mathcal{F}_d \; \forall d \in \mathcal{D}$.
- Limiting the Combat Elements (CEs) $\mathcal{C}$ to be supplied to $\mathcal{C}_d \; \forall d \in \mathcal{D}$.

Depending on the coordination process, coordination could range from completely disjoint subsets (no cooperation possible), to complete subsets, such that the problem equals the fully connected case in Figure 4-1. Proving that coordination increases scalability and robustness

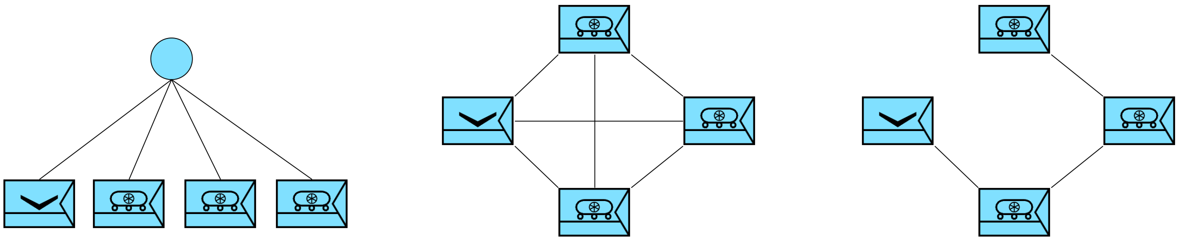

Figure 4-1: The different subset approaches are shown. Left, a central node solves the problem for all vehicles. In the middle, the the subsets are complete, such that all vehicles cooperate with each other on all variables. At the right, each vehicle receives a distinct but overlapping subset, such that each agent can cooperate with a specific set of neighbors.

against communication failures, is out of scope for this thesis. Instead, ensuring that the solution method allows coordination is assumed sufficient. For simplicity, the subsets are defined to be complete, such that the topology equals the central case displayed in Figure 4-1, and all vehicles are neighbors with all other vehicles.

4-1-2 Cooperation

After coordination, each vehicle can cooperate with its neighbors. Cooperation can then result in vehicles synchronizing their actions to perform a physical transfer, or by communicating their actions, such that their decisions can be aligned to pursue global optimality. Decisions that require cooperation by synchronization can be described as:

- Transferring $T_{kdd't}$ supplies of class $k \in K$ from vehicle d' to $d \in \mathcal{D}$ at time $t \in T$.
- Refueling amount $R_{dd't}$ from vehicle d' to $d \in \mathcal{D}$ at time $t \in T$.

The act of receiving the supplies or fuel is denoted as a collection. Vice-versa, the providing vehicle agrees with a constraint, which is the opposite obligation to satisfy the collection by another vehicle.

As stated, the vehicles also communicate their local decisions, including delivery D_{kdct}, pickup P_{kdt}, and relocation x_{dijt}. These decisions are coupled implicitly through the inventory capacity at the CEs b_{kc}, the inventory I_{kct}, thresholds τ_{kct} and weights w_{kct}, such that globally a higher objective value can be achieved by aligning their decisions.

To enable distributed cooperation, two algorithms are proposed and adjusted for synchronization, based on the Distributed Constraint Optimization Problem (DCOP) paradigm. The first is Maximum Gain Messaging (MGM), proposed by Maheswaran, Pearce, and Tambe (2004). MGM is an incomplete, synchronous algorithm performing distributed local search. It is also an any-time algorithm, meaning that over time the solution can only improve. As an extension for MGM, MGM-2 is proposed (Maheswaran et al., 2004). The essence of this extension is that agents performing the algorithm can explicitly propose a collection or constraint, as opposed to MGM. Both algorithms are explained further in Section 4-2-2.

4-1-3 Local problem

Independent of any prior coordination or cooperation being possible, each vehicle always solves its local problem. In the case of no cooperation, individual solutions are still generated, though their subsets might overlap. Global feasibility remains ensured, as even if vehicles try to supply a CE already supplied by another vehicle, the inventory or fuel levels of either the CEs or vehicles will never be lower than expected. Hence, vehicles will always be able to perform remaining deliveries or return to an FOB respectively. When cooperation is possible, vehicles can claim any excess supplies or fuel of a neighboring vehicle, resulting in a collection and corresponding constraint.

To solve the local problem, a two-stage heuristic is proposed. During the first stage, a greedy delivery scheduling algorithm proposes the deliveries, pickups, transfers, and relocations. In the second stage, a path planning algorithms creates a path that avoids as much risk as possible between the scheduled decisions. The method is based on the exact forward-labeling algorithms for the Elementary Shortest Path Problem with Resource Constraints (ESPPRC) (Feillet, Dejax, Gendreau, & Gueguen, 2004) (Bettinelli, Ceselli, & Righini, 2011) (Pugliese & Guerriero, 2013). In general, this method functions by dynamically extending profitable labels, and removing dominated labels which will certainly produce a worse outcome than at least one other label. In Section 4-2 the delivery scheduling and path planning algorithm are explained further.

4-2 Local algorithms

4-2-1 Delivery scheduling

Decisions

During the delivery scheduling, each vehicle $d \in \mathcal{D}$ generates its main decisions. The decisions that can be made in each period $t \in T$ can be summarized as:

- Delivering an amount of D_{kdct} supplies of each class $k \in K$ to CE $c \in \mathcal{C}$.
- Resupplying P_{kdt} at any FOB up to the capacity b_{kd} of each class $k \in K$.
- Collecting $T_{kdd't}$ supplies up to b_{kd} of each class $k \in K$ from vehicle $d' \in \mathcal{D}$.
- Collecting $R_{dd't}$ fuel from vehicle $d' \in \mathcal{D}$, or recharge R_{dt} up to r_d at any FOB.
- Relocate from node i to node $j \in N$ at time $t \in T$.
- Satisfying constraints, by transferring $-T_{kd'dt}$ or recharging $-R_{d'dt}$ to vehicle $d' \in \mathcal{D}$.

To be able to deal with the constraints, sufficient fuel or supplies must be kept at hand. To keep track of these amounts, the set of constraints for each vehicle is denoted as $\mathcal{H}_d$. Then, during planning, the required amounts of supplies or fuel that need to be present at time $t \in T$ for constraint satisfaction are defined as m_{dkt} $\forall k \in K$ and m^e_{dt} respectively.

As the objective is to minimize the maximum penalty, a possible delivery to a CE is calculated using the following formula:

$$D_{kdct} = \min(\max_{t \dots T}(\tau_{kct} - T_{kct})^+,\ I_{kdt} - m_{dkt},\ b_{kc} - I_{kct}) \quad \forall k \in K,\ d \in \mathcal{D},\ c \in \mathcal{C} \tag{4-1}$$

This formula states that first, at most the maximum shortage for the CE over all future periods is delivered, limited by the available inventory of the vehicle, from which any possible supply margins m_{dkt} are subtracted. Then, this amount is limited further by the available inventory space at the CE.

Delivery selection

To prevent vehicles from deciding to do a profitable, but unacceptable risky delivery, the deliveries are discounted by an approximated risk. Because the precise path planning is not yet performed, a single-hop approximation is implemented. It is assumed that with one intermediate point, a relatively safe path can be constructed. Time limits are not included in this approximation, such that very long, but safe paths might be accepted as a valid approximation, while not being possible. The benefit of this relaxation is that the approximation for each pair of nodes can be precomputed as such:

$$\hat{\mathbb{P}}_{dijt} = 1 - \max_{\forall s \in N}\{(1 - \mathbb{P}_{dist}) \cdot (1 - \mathbb{P}_{dsjt})\} \quad \forall d \in \mathcal{D},\ t \in T \tag{4-2}$$

Based on this approximation, the discounted delivery performed at time $t' \in T$ is defined as $\hat{D}_{kdct'} = D_{kdct'} \cdot \hat{\mathbb{P}}_{dijt}$, with $i \in N$ being the current position of vehicle $d \in \mathcal{D}$ at time $t \in T$, and $j \in N$ the position of CE $c \in \mathcal{C}$ at time $t' \in T$. Subsequently, this discounted delivery can be used to calculate an approximation of the expected CE inventory and the corresponding maximum weighed penalty, using Equation and (3-3) Equation (3-4) respectively.

For each possible delivery, the vehicle can then calculate the gain based on the approximation of the expected penalty. The delivery selected is the one maximizing the gain. As it might occur that the maximum penalty cannot be reduced, while the delivery is still beneficial, a hierarchical gain is implemented:

1. Reduce the maximum penalty of all CEs combined.

2. Reduce the maximum penalty by the largest amount of any CE.

3. Reduce as much penalty as possible.

The examples below illustrate how this hierarchical selection functions. In example 1, the maximum penalty of 90 of all CEs can be reduced by delivering to CE 1. In example 2, the overall maximum penalty cannot be reduced, but delivering to CE 2 would reduce at least that maximum penalty. In example 3, neither of the maximum penalties can be reduced, but the total penalty can be reduced by delivering to CE 1.

	Example 1		Example 2		Example 3	
Delivery to	CE 1	CE 2	CE 1	CE 2	CE 1	CE 2
Max penalty before	90	70	90	70	90	70
Max penalty after	65	15	90	15	90	70
Total penalty before	800	900	800	900	800	900
Total penalty after	400	600	400	600	400	600

Calculating fuel & supply margins

A constraint essentially means that during delivery scheduling of vehicle $d \in \mathcal{D}$, the resources that are collected by vehicles $d' \in \mathcal{D}$ must be made available in the agreed period and node. To plan the required amounts at time $t \in T$, the supply margins m_{dkt} $\forall k \in K$, and the fuel margins as m^e_{dt} have to be calculated. This section explains how the margins are calculated based on the available set of constraints $\mathcal{H}_d$.

The main influence on the required margins is whether an FOB is reachable before the constraint needs to be satisfied. If an FOB is not reachable before the constraint, sufficient fuel and supplies need to be collected and preserved earlier. To verify this, a short algorithm is executed over the constraints:

1. By iterating forward, retrieve the path maximizing the fuel level at the next constraint, either by going directly, via an FOB or a collection. However, if an FOB is forced, the direct path is forbidden (see step 3).

2. By iterating backward, set the current required inventory and fuel margins, such that if no FOB is reachable not only the current constraint values are required, but also of all succeeding constraints that cannot be reached by an FOB prior.

3. By iterating forward, passing via an FOB is forced for any case in which a direct path is chosen to maximize fuel, but there is not sufficient inventory to fulfill the supply requirement.

These steps are repeated iteratively until no more new FOBs are forced to be visited. If during the first step multiple collections can be visited before the next constraint, it is assumed that they are visited either all or none. To clarify why the above steps are necessary, an example is provided in Figure 4-2.

Recharging enables FOB: in example 4-2a, the vehicle has a fuel capacity of 6 and it starts with a fuel level of 4. In this case, the FOB between constraints A and B can only be passed if the FOB is passed before constraint A too. In other words, maximizing the fuel level upon arriving at A, enables an FOB later. Though visiting the FOB before might initially be unnecessary to fulfill constraint A, it can then still be forced such that later constraints can be satisfied.

Resupplying prohibits FOB: in example 4-2b, the vehicle starts with a fuel level of 5 instead. If in this case the FOB is visited before constraint A, there is not enough fuel remaining to visit the FOB between constraint A and B, meaning that it needs to collect sufficient supplies before A for all subsequent constraints.

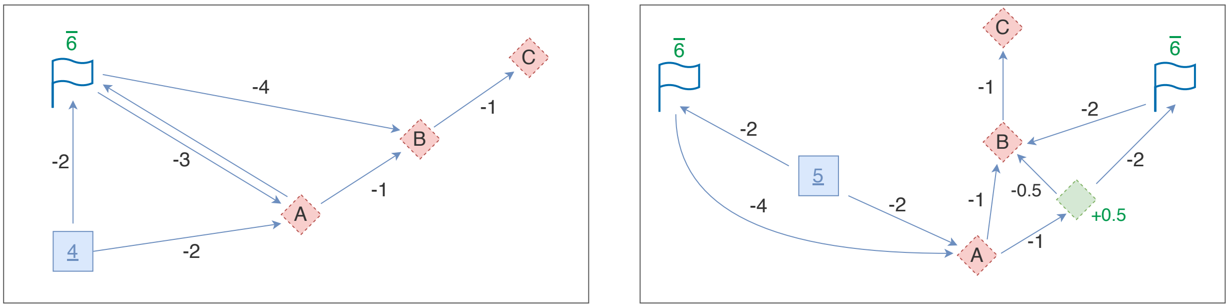

(a) **(b)**

Figure 4-2: The blue square depicts the vehicle with the current fuel level, the arcs indicate the fuel consumption, the red diamonds the constraints, the flags with the fuel vehicle capacity in green represent FOBs, and the green diamonds the fuel collections.

Algorithm overview

To aid the solution process, the scheduling is split in several cases, which are all covered in the pseudo-code from Algorithm 1 displayed below.

Algorithm 1: Greedy delivery planning for vehicle $d \in \mathcal{D}$

while *any decision is possible* **do**
    if *any constraints* **then** track minimum required fuel and supply levels
    if *any vehicle inventory class empty* **then** require resupply
    if *no resupply or refueling required* **then**
        collect all reachable CE positions
    else
        collect all reachable CE positions via FOBs or neighbors

    foreach *reachable CE position* **do**
        collect maximum penalty and total penalty reduction
    select most profitable reduction as next delivery

    if *no profitable delivery possible* **then**
        if *inventory or fuel level below constraint levels* **then** require resupply or refueling
        if *resupply or refueling required* **then** go to FOB or make collection if possible
        if *no resupply or refueling required or impossible* **then**
            case *possible before constraint* **do** go to rebalancing node
            case *possible before constraint* **do** return to current node
            otherwise do go to constraint node

    if *fuel threshold not violated* **then**
        update position, time, fuel and inventory levels
    else
        discard delivery and require refueling

return *delivery plan consisting of deliveries and relocation decisions*

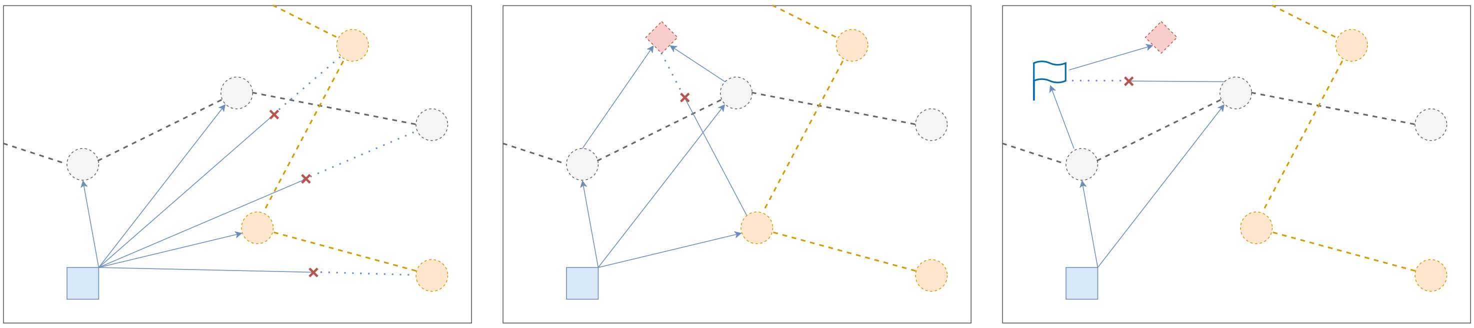

(a) Three positions are directly reachable **(b)** Two are reachable before a constraint **(c)** One before a constraint via an FOB

Figure 4-3: Three subcases for direct delivery are displayed. The orange and grey circles indicate the future positions of two CEs. The blue square indicates the current position of the vehicle. The red diamond designates a constraint, meaning a transfer agreement with another vehicle, and the blue flag depicts an FOB.

Case 1: Direct delivery

The simplest case of delivery scheduling is a direct delivery from vehicle $d \in \mathcal{D}$ to any CE $c \in \mathcal{C}$, before which no resupply or refueling is necessary. In Figure 4-3 the choices for a direct delivery are depicted. There are three subcases. The first, depicted by Figure 4-3a, is that of a vehicle being able to reach specific future positions of a CE both in time and with its current fuel level. In Figure 4-3b, also a constraint must be satisfied after delivery, limiting the possibilities. Thirdly, in Figure 4-3c the constraint can be reached after delivery, but it is not always possible to visit an FOB before. If it is possible, the supply margins m_{kdt} are set to zero, and the fuel margins m^e_{dt} to the fuel needed to reach the FOB.

Case 2: Delivery after resupply or refueling

In the second case, either resupply or refueling is required. Resupply is automatically required when a vehicle $d \in \mathcal{D}$ has run out of any inventory class at time t, meaning $\exists\, k \in K : I_{kdt} = 0$. Refueling is required if the previous attempt to schedule a delivery resulted in a proposed node $i \in N$ that was out of reach from the closest FOB at node $j \in N$, meaning $E_{dt} - C_{dij} < 0$. Furthermore, resupply is also required if the current inventory at time t is lower than the margin, meaning $\exists\, k \in K : I_{kdt} - m_{dhk} < 0$, or if the current fuel level at node $i \in N$ is not sufficient to reach and satisfy the constraint at node $j \in N$, meaning $E_{dt} - C_{dij} - m^e_{dh} < 0$.

Given that resupply or refueling is required, the two subcases are depicted in Figure 4-4. First, an FOB could be visited before a delivery, limiting the reachable future positions, as displayed in Figure 4-4a. In the second case, there are future positions that are not reachable via an FOB, but are reachable via another vehicle, allowing for a fuel and supply transfer (Figure 4-4b). In both cases, however, any constraints still need to be met. Therefore, also the subcases described in Figure 4-3b and 4-3c apply.

It is possible that refueling is required, and that a scheduled delivery after refueling is again outside of reach of the closest FOB. This would trigger refueling again, resulting in a loop. To prevent this, those deliveries are allowed if more than 25% of the fuel capacity is left, which is assumed sufficient to make a near collection. If it is lower, the delivery is discarded.

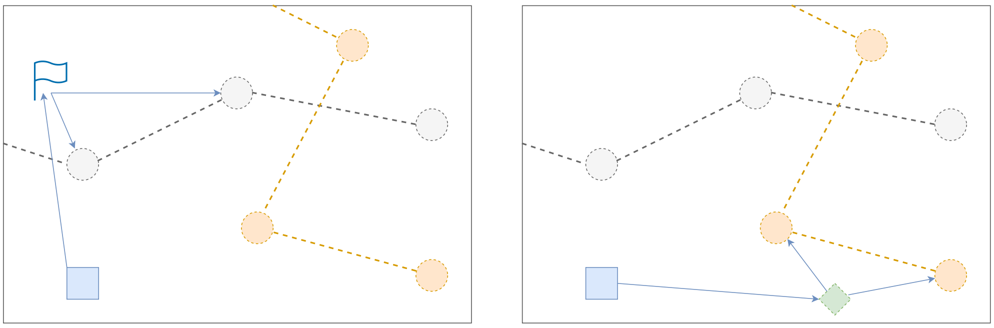

(a) Two positions are reachable via an FOB **(b)** Two are reachable via another vehicle

Figure 4-4: Two subcases are displayed for delivery after refueling or recharging. The circles depict future CE positions, the blue square the current vehicle's position, the blue flag an FOB, and the green diamond designates a vehicle available for transfers.

Case 3: Direct resupply or refueling

The third case emerges when a vehicle is required to recharge or resupply, but no delivery is possible or profitable. If possible, the closest FOB is selected, or otherwise the closest vehicle that has any excess resources. However, this is only allowed if after the transfer the supply margin m_{dhk} and fuel margin m^e_{hk} are satisfied, or if the available fuel transfer allows reaching an FOB that is otherwise not reachable. Again, if constraints are present, the case depicted in Figure 4-3b still applies. Regardless of refueling being successful or not, it is not checked whether the remaining fuel level is sufficient to reach the closest FOB. This means that the refueling flag cannot be triggered again, which could otherwise result in a loop, as the nearest refueling action, if possible, has already been scheduled.

Case 4: Relocation

Relocation covers the three remaining subcases. First, the vehicle can rebalance to a new position, possibly better suiting the planning of other vehicles and enabling more cooperation. During rebalancing, the destination is selected as the node minimizing the summed distance to all other vehicles and CEs, weighed by the one-hop transit risk for going from the current node $i \in N$, to the destination node $j \in N$ at time $t \in T$:

$$\underset{\forall j \neq i \in N}{\arg\min} \left\{ \hat{\mathbb{P}}_{dijt} \cdot \left(\sum_{\forall s \in S_d} d_{js} + \sum_{\forall s \in S_c} d_{js} \right) \right\} \tag{4-3}$$

The nodes occupied by neighbors are defined as $S_d = \{j \in N : y_{d'jt} = 1 \ \forall d' \neq d \in \mathcal{D}\}$, the CE nodes as $S_c = \{j \in N : y_{cjt} = 1 \ \forall c \in \mathcal{C}\}$, and the distance between two nodes $i, j \in N$ as d_{ij}. Possibly, the selected node is further away than allowed. In that case, the node that is closest to the selected destination node, but is still reachable, is set as the new destination.

Alternatively, if due to constraints the rebalancing node is not possible, the vehicle can decide return to its current position after a fixed time, allowing the possibility for constructing a safe path near the current position. In this case, the fixed time is set to 2 hours. Lastly, the vehicle can go towards the meeting node belonging to the next constraint.

4-2-2 Path planning

During the second stage, the objective is to create a path minimizing the total risk taken, while performing the scheduled actions. During path planning, three decisions can be made:

- Wait a full period at a node.
- Relocate to a new node.
- Resupply or refuel if an FOB is located at the node, or if a scheduled collection is available at the considered node and period. This decision can be taken in parallel with the first two.

Though this approach is based on an exact method, this specific method loses that property to speed up the solution process, as labels are initially only extended to the three nodes closest to the current node, the current node itself, the meeting node, all FOB nodes, and any scheduled transfer node. By doing this, the search space is severely reduced while not losing any feasibility guarantees. As soon as a feasible solution has been found, the search is reinitialized, and all nodes all enabled for expansion. The pseudo-code is provided in Algorithm 2.

Algorithm 2: Dynamic path planning for vehicle $d \in \mathcal{D}$

```
foreach action in the delivery plan do
    initialize source label with current position, fuel and inventory levels
    if currently at FOB or scheduled transfer available then resupply and/or refuel

    while not terminated and labels left do
        select new label v_l to expand
        foreach node do
            expand label v_l to v'_l
            if label infeasible in time or fuel then discard label
            if FOB at node or scheduled transfer available then resupply and/or refuel
            if insufficient fuel or supplies and (all FOBs out of reach or scheduled
              transfer already performed or scheduled transfer out of reach) then discard
              label
            if node is meeting node and arrival time in meeting period then
                check if new best path and if so, update
            else
                discard label
        if highest remaining survival probability lower than best found or no improvement
          found within time limit then terminate preemptively
    recursively build path segment from current best label
    remove any excess fuel or supplies from the scheduled transfer
add path segment to complete path

return complete path and intermediate resupply and refueling transfers
```

Extending paths

First, a label $v_l = \langle i, \mathbb{P}^s, \hat{\mathbb{P}}^s, t^a, E, I_k, \mathrm{f}, n, l' \rangle$ is introduced, tracking the current node i, the cumulative current $\mathbb{P}^s$ and expected $\hat{\mathbb{P}}^s$ probability of survival, the arrival time t^a, the current energy and inventory level E and I_k respectively, a Boolean f indicating if the scheduled transfer has already been performed, the number of consecutive periods n it has remained static at the current node, and the preceding label index l'.

The label v_l with the highest expected probability of survival $\hat{\mathbb{P}}^s$ is selected for extension, such that effectively a depth-first search is implemented. This expected probability of survival for vehicle $d \in \mathcal{D}$, when planning to the next scheduled action at node $j \in N$ in period p, is approximated using the transit risk once, and the static risk over the remaining periods:

$$v_l\langle \hat{\mathbb{P}}^s \rangle = v_l\langle \mathbb{P}^s \rangle \cdot \left(1 - \hat{\mathbb{P}}_{d, v_l\langle i \rangle, j, v_l\lceil \langle t^a \rangle \rceil}\right) \cdot \left(1 - \mathbb{P}^0_{d,j,\lceil \langle t^a \rangle \rceil}\right)^{p - \lceil \langle t^a \rangle \rceil - t_{v_l\langle i \rangle, j}} \tag{4-4}$$

the new label v_l is then extended to a selection of, or all nodes $j \in N$. When extending, the performed updates can be described as:

$$t_a = \begin{cases} \lceil v_l\langle t^a \rangle \rceil + 1 & \text{if } v_l\langle i \rangle = j \\ v_l\langle t^a \rangle + t_{v_l\langle i \rangle, j} & \text{if } v_l\langle i \rangle \neq j \end{cases} \tag{4-5}$$

$$n = \begin{cases} v_l\langle n \rangle + 1 & \text{if } v_l\langle i \rangle = j \\ 1 & \text{if } v_l\langle i \rangle \neq j \end{cases} \tag{4-6}$$

$$E = \begin{cases} r_d & \text{if } \sum_{f \in \mathcal{F}} y_{fj} \geq 1 \\ v_l\langle E \rangle - C_{v_l\langle i \rangle j} + \sum_{d' \in \mathcal{D}} R_{dd'\lceil t^a \rceil} & \text{if } \sum_{d' \in \mathcal{D}} y_{dj\lceil t^a \rceil} \geq 1 \\ v_l\langle E \rangle - C_{v_l\langle i \rangle j} & \text{if } v_l\langle i \rangle \neq j \\ v_l\langle E \rangle & \text{Otherwise} \end{cases} \tag{4-7}$$

$$I_k = \begin{cases} b_{kd} & \text{if } \sum_{f \in \mathcal{F}} y_{fj} \geq 1 \\ v_l\langle I_k \rangle + \sum_{d' \in \mathcal{D}} T_{kdd'\lceil t^a \rceil} & \text{if } \sum_{d' \in \mathcal{D}} y_{dj\lceil t^a \rceil} \geq 1 \wedge \mathrm{f} = 0 \\ v_l\langle I_k \rangle & \text{Otherwise} \end{cases} \tag{4-8}$$

$$\mathrm{f} = \begin{cases} 1 & \text{if } \sum_{d' \in \mathcal{D}} T_{kdd'\lceil t^a \rceil} > 0 \vee \sum_{d' \in \mathcal{D}} R_{dd'\lceil t^a \rceil} > 0 \\ v_l\langle \mathrm{f} \rangle & \text{Otherwise} \end{cases} \tag{4-9}$$

$$\mathbb{P}^s = \begin{cases} v_l\langle \mathbb{P}^s \rangle \cdot (1 - \mathbb{P}_{dj\lceil t^a \rceil}) & \text{if } v_l\langle i \rangle = j \\ v_l\langle \mathbb{P}^s \rangle \cdot (1 - \mathbb{P}_{dv_l\langle i \rangle j\lceil t^a \rceil}) & \text{if } v_l\langle i \rangle \neq j \end{cases} \tag{4-10}$$

In Equation 4-5, the arrival time is updated. This can either be the result of static waiting or relocation to a new node. Equation 4-6 updates the number of consecutive periods based on the same options. In Equation 4-7, the fuel level is increased to the vehicle capacity at an FOB, decreased with the consumption during relocation, kept constant during waiting, or increased with any scheduled recharging. A similar case is made for the inventory level in Equation 4-8. If the collection has been performed, the transfer Boolean is updated in Equation 4-9. Lastly, the survival probability is updated using Equation 4-10, in which either the probability for a strike for remaining static or being in transit is used. Also the expected survival probability $\hat{\mathbb{P}}^s$ needs to be updated using Equation (4-4).

Removing labels

First, some general rules can be formulated that prove that labels cannot result in a feasible or optimal solution. In those cases, they can be removed. The meeting period is defined as p, the current and meeting node as i and $j \in N$ respectively, and the required margins as m_{dhk} and m^e_{dh} for fuel and inventory. The rules are formulated as follows:

- The survival probability is lower than the currently best known, meaning $\mathbb{P}^s < \mathbb{P}^s_best$
- The node can never reach the scheduled action on time, as $t^a + t_{ij} > p$.
- Relocation cannot be performed because of lacking fuel, as $E - C_{ij} < 0$.
- The margins cannot be met anymore, as no FOB or scheduled collection can be reached, meaning $I_{dkt} < m_{dkt} \vee E_{dt} < m^e_{dt}$.

As the probability for survival monotonically decreasing, it can never increase, meaning it directly serves as an upper bound on the best solution that can be expected by extending an existing label, allowing the use of the first rule. This also means that, if the highest survival probability of all available labels is lower than the best-found probability so far, the search can be terminated, as the best solution has been found.

Unfortunately, the dominance criteria used for the ESPPRC cannot be applied directly to this problem. The main difference is that in this case, the triangle inequality does not hold. For example, in Figure 4-5a a longer, safer path exists. If sufficient time is available, it yields a better solution. In Figure 4-5b no such path exists, such that more time will only lead to longer waiting, and thus a worse solution. Therefore dominance of label v_l over v'_l is only exact when they occupy equal nodes at the same time, with label v_l having equal or larger fuel and inventory, is present equal or less consecutive periods, and has a higher survival probability. This kind of dominance is very weak, for which the computing power for comparison can easily exceed the profit of reducing the search space slightly.

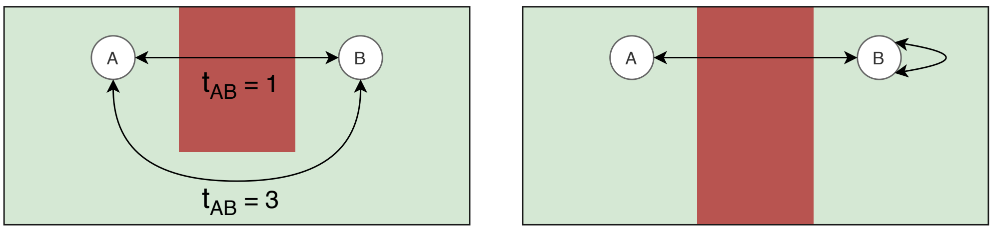

(a) A longer safer path exists **(b)** No safer path exists

Figure 4-5: An example showing that a longer time left can result in both a safer or riskier path

To compensate this, the heuristically reduced search space along with the approximation of the expected survival probability is used. Still, due to the exponential complexity search time can easily explode. To prevent this, all actions during greedy scheduling are limited to the number of periods it takes to cover a distance of 80km, limiting the worst-case number of possible paths independent of the vehicle's speed.

4-3 Cooperative algorithms

4-3-1 Maximum Gain Messaging

Calculating excess fuel & supplies

Before any cooperation can take place, the excess supplies that can be collected have to be calculated. After a vehicle $d \in \mathcal{D}$ has completed its delivery schedule and precise path planning, it can calculate the excess supplies for time $t \in T$, resulting in a value $g_{dkt}\ \forall k \in K$, and any excess fuel as g^e_{dt}. These values are then communicated during cooperation, such that any neighboring vehicle can collect these excess resources, resulting in a constraint.

The amount of excess resources are explained using the simplified example below. For each period t, a figurative inventory level I_t, delivery D_t and excess inventory g_t are displayed. In period 4, the inventory is increased up to the capacity of 10 at an FOB. As can be seen, before visiting the FOB, the vehicle has 5 supplies in excess. During and after resupply, however, it has none in excess as it has to deliver its complete inventory.

t	1	2	3	4	5
I_t	10	5	5	10	0
D_t	0	5	0	0	10
g_t	5	5	5	0	0

In reality, the excess supplies are also reduced by planned constraints and increased by collections. Furthermore, instead of deliveries the remaining fuel consumption for executing the planned path is taken for the fuel margin. At the last period, also the fuel consumption to reach the nearest FOB is deducted, such that no vehicle can claim more fuel than is necessary for the other vehicle to return to an FOB.

Solution process

MGM executes in synchronized rounds. During each round, all vehicles optimize their new plan locally, based on the most recent information. After creating a new plan, all vehicles broadcast the gain achieved with their new plan. If a vehicle has the highest gain of all its neighbors, it can directly accept the new plan. If not, it discards its new plan for this round. The steps of the algorithm are depicted in Figure 4-6.

The gain for selecting the best new plan is calculated similarly as when picking the next delivery during greedy delivery scheduling. The main difference is that the decision can also be based on the average survival probability. The order then becomes:

1. Largest reduction of the maximum penalty of all CEs combined.
2. Largest reduction of a maximum penalty of any CE.
3. Largest reduction of total penalty.
4. Largest increase in average survival probability over all neighbors.

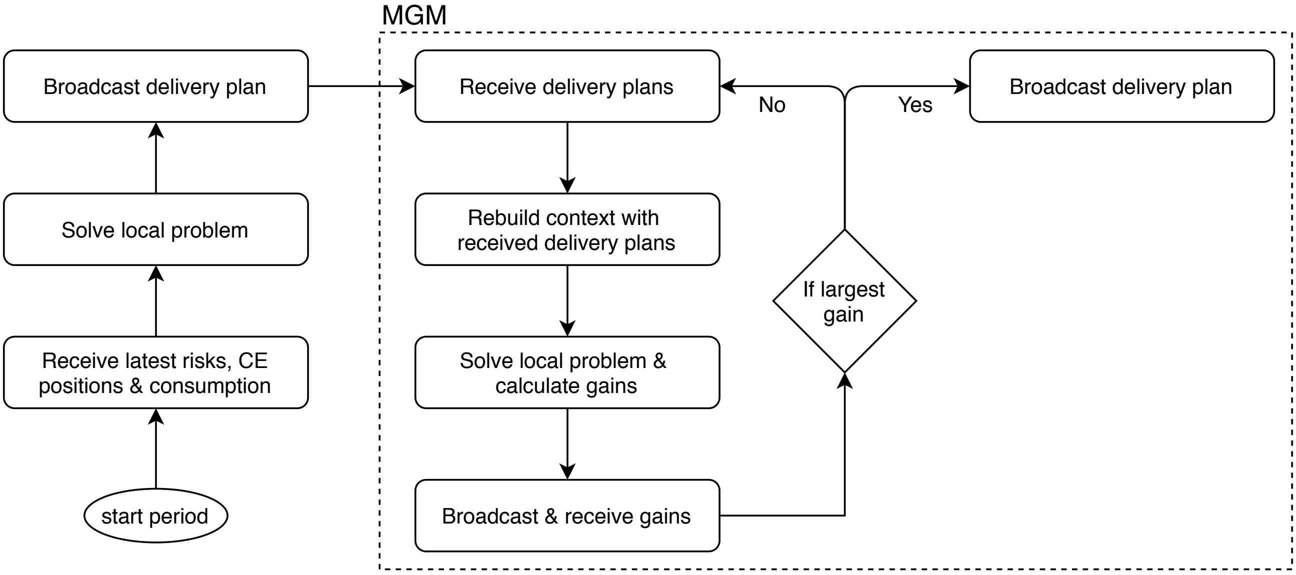

Figure 4-6: A flow chart depicting the different steps of the MGM algorithm. Initially, each vehicle solves the problem individually. Then rounds of MGM are performed to allow cooperation.

The flow chart in Figure 4-6 shows the solution process. At the start of each period, all vehicles receive the most recent information. Then, they individually solve their local problem. This solution is already a feasible, non-cooperative solution. With these plans, all vehicles participate in rounds of MGM, until no more improvements are found or if the maximum number of rounds is reached.

The step 'rebuilding the context with the received delivery plans' means that the vehicle includes the deliveries of other vehicles in the future shortages of CEs, and can now see if other vehicles have any fuel or supplies in excess, meaning it can subsequently plan to collect those. Each time a delivery plan is accepted including collections, the corresponding constraints are created at the providing vehicle. To prevent infeasibility, the providing vehicles can never cancel a constraint using MGM, but the receiver can unilaterally do so.

If a vehicle has created a new delivery plan that resulted in the largest gain of its neighbors, will broadcast it. In the next round it does not have join, as it has not received any new information and as such cannot improve over its previous plan. In the round after, it will join again.

4-3-2 Maximum Gain Messaging 2

Only using MGM would limit the flexibility during scheduling severely. The problem is that with MGM, vehicles cannot remove any of their constraints. Furthermore, the only available resources are excess resources, which is not necessarily optimal. Thus, the importance of adding MGM-2 is that it enables vehicles to alter their synchronized moves. Put otherwise, providing vehicles can propose to cancel a constraint. Furthermore, it allows vehicles to generate proposals for yet unavailable fuel or supplies, effectively proposing to another vehicle to free up those resources.

During each round of optimization, vehicles are randomly assigned to become an offerer or a receiver. If the offerer has any constraints, it will propose to remove one. If it has no constraints it will try to propose a collection. If two updated delivery plans are both feasible and the overall gain is positive, they are accepted and broadcasted.

Proposing to remove a constraint

First, the offerer randomly selects any of its constraints to remove. The linked vehicle then automatically becomes the receiver. When removing a constraint, the offerer calculates its new delivery plan without the constraint first, while it is not allowed to schedule a collection from the receiving vehicle, as the currently known excess resources of the receiver vehicle are possibly dependent on the constraint to be removed. After the offerer has revised its plan, it sends it to the receiver along with the gain. At the receiver side, infeasibilities might still occur, as the collection could be necessary to satisfy other constraints of the receiver. Hence, removing its corresponding collection would not be possible. On successful rescheduling of the receivers' delivery plan, the plans are both accepted if the total gain is positive.

Proposing to add a constraint

If the offerer does not have a constraint, it will check if it is valuable to propose one. It can either propose to collect fuel or supplies, but not both. The vehicle will try to collect fuel only if its current fuel level is below 25% of its capacity, or when it is out of reach of the nearest FOB. Otherwise, it will try to collect supplies.

In both cases, the offerer selects the receiver based on a normalized distance-availability metric, such that the most fuel available is favored, as well as the closest distance:

$$\underset{d' \neq d \in \mathcal{D},\ j' : y_{d'j't}=1}{\arg\min} \left\{ \frac{\max_{d'' \neq d \in \mathcal{D}} \{E_{d''t}\}}{E_{d't}} \cdot \frac{\min_{j'' \in N : y_{d''jt}=1} \{d''_{ij}\}}{d_{ij}} \right\} \tag{4-11}$$

For supplies the same equation is applied, but instead of the available fuel E_{dt}, the sum of all classes $\sum_{k \in K} I_{kdt}$ is used.

When a receiver is selected, the maximum obtainable fuel is defined as 35% of the receiver's fuel level, or the maximum available at the offerer. For inventory, this limit is increased to 50% and any existing supply constraints of the receiver are subtracted. Then, only if the receiver can reach the position of the offerer and any fuel or supplies can be obtained, the collection is proposed. The transfer node is set as the current node of the offerer, and the period as the latest arrival time of either vehicle at this node. Subsequently, the constraint is added to the receiver, and the collection is added as a first action during the greedy scheduling of the offerer.

When an attempt is made to solve the local problems with this new proposition, the receiver plans first, and may not cooperate with the offerer as any information on excess resources is still obsolete. After completion, the offerer reschedules too and checks if a net gain is present. If so, both new plans are accepted.

Chapter 5

Data

To be able to simulate the effectiveness of the proposed solution method for Dispersed Autonomous Resupply (DARE), a hypothetical scenario on familiar territory against a peer opponent is considered. Using this scenario, realistic data can be generated.

5-1 Scenario

Unexpectedly, the UK attacks Belgium and the Netherlands. The aim is to secure both the capitals Amsterdam and Brussels, and the main ports in Rotterdam and Antwerpen, such that it can maintain itself as the designated power in North-West Europe. Following a swift landing by the Royal Marines near Antwerpen, a firm beachhead has been established, such that the UK's 12^{th} Armored Infantry Brigade can move towards Rotterdam. The Dutch 13^{th} Light Brigade is tasked to intercept and halt the 12^{th} Armoured Infantry Brigade before Rotterdam is reached, such that sufficient time can be created for mechanized reinforcements to arrive, which is expected to take 72 hours.

As depicted in Figure 5-1a, the 12^{th} Armoured Infantry Brigade is expected to advance over two axes as quickly as possible to capture Rotterdam. The 12^{th} consists of a light tracked reconnaissance battalion, a light mechanized battalion with Armoured Personnel Carriers (APCs) and Light Armoured Vehicles (LAVs), two heavy tracked mechanized battalions with Infantry Fighting Vehicles (IFVs), and one armoured battalion with tanks. The light component is expected to advance as fast as possible over the western flank via Bergen op Zoom, to take the Hellegatsplein crossing and the Haringvliet bridge of the A29 highway. The heavy component is expected to advance over the eastern flank via Breda, to take the railway and A16 highway Moerdijk bridges. The own forces consist of two wheeled, mechanized infantry battalions with APCs, one wheeled, mechanized combat engineer battalion, and one reconnaissance squadron.

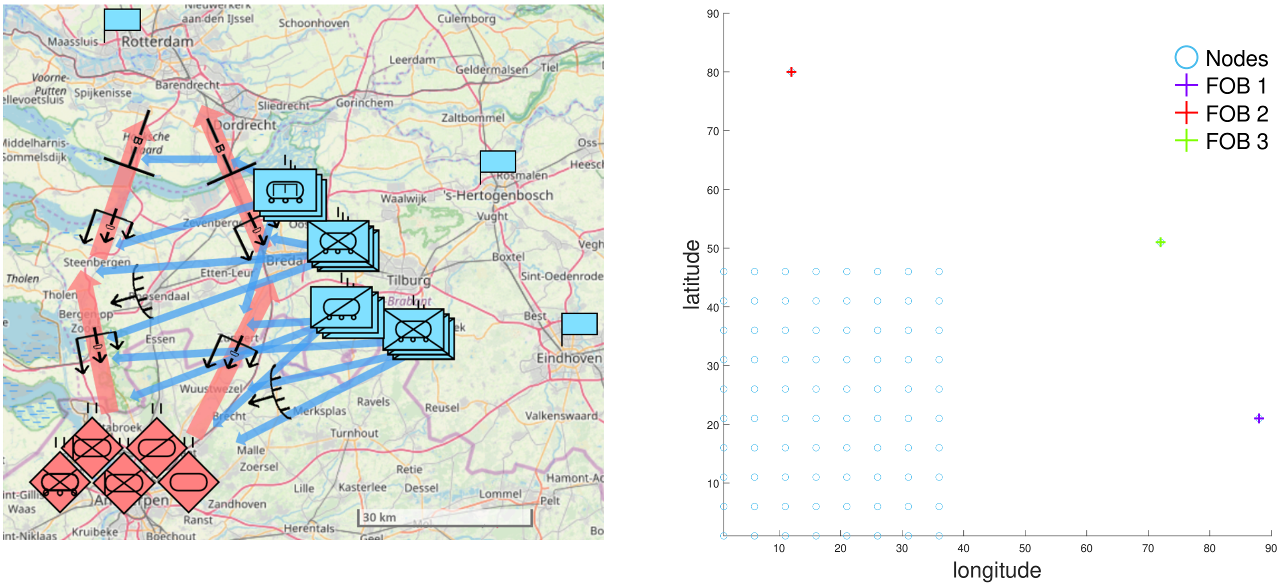

(a) The Concept of Operations (CONOPS).

(b) The sampled Area Of Operations (AOO).

Figure 5-1: The hypothetical CONOPS is shown with the corresponding sampled version of the environment, by creating equally spaced nodes at 5 kilometers distance within the AOO, with three Forward Operating Bases (FOBs) at approximately 90 kilometers distance.

5-2 Quantifying operations

The described scenario is used to create an approximate representation of distances, movements, risks, and supply consumption. This is no one size fits all, but it serves as an approximation for dispersed operations against a peer opponent.

Environment

First, a sampled environment is generated. As no large rivers, mountains, swamps or dense forests are present, all vehicles are assumed to be able to travel to any location. The operation sketch from Figure 5-1a provides a rectangle area of operations of around 40km by 50km. The area is sampled with equally spaced nodes every 5km, to limit the problem size. Three FOBs are placed roughly 90km away from Antwerpen, at Eindhoven, 's Hertogenbosch, and Rotterdam. The sampled environment is depicted in Figure 5-2b.

Enemy movement

Each enemy is assigned to an axis of advance, as depicted in Figure 5-1a. The paths are sampled such that in 72 hours the target is reached. The expected paths are not exact, as uncertainty is introduced by randomly perturbing the assigned paths. Each location $\{l_y, l_x\}$ along the path is perturbed using a normal distribution, to create the new locations $l'_y = \mathcal{N}(l_y, 0.04)$ and $l'_x = \mathcal{N}(l_x, 1.44)$, such that a standard deviation of 0.2 and 1.2 kilometers is present in the latitudinal and longitudinal direction respectively.

Table 5-1: The total supply capacity and fuel consumption of one platoon for all classes per type.

Type	Class I	Class II	Class III	Class IV	Class V	Total	consumption
Infantry	600kg	150kg	2.400kg	200kg	3.200kg	6.550kg	2.8 L/km
Engineers	600kg	320kg	2.500kg	350kg	2.800kg	6.570kg	3.0 L/km
Reconnaissance	260kg	140kg	1.600kg	190kg	1.700kg	3.890kg	1.8 L/km

CE movement

First, the enemy forces are uniformly assigned to the Combat Elements (CEs). Furthermore, for each CE a center point is generated by uniformly distributing them over the AOO. To simulate the movement of the CE, nodes are selected to create a path. Assuming a maximum speed of 40 km/h, the set of possible nodes N_{lim} is limited to those it can reach within one period. The nodes $j \in N_{lim}$ are ordered based on the following distance measure:

$$\hat{d_j} = d_{ij} + d_{j,ce} + 35 \cdot \left(\frac{d_{j,cp}}{35}\right)^2 \tag{5-1}$$

Here, $i \in N$ is the current CE node, $d_{j,ce}$ describes the distance to the closest enemy, and $d_{j,cp}$ to the center point. The latter is weighed and squared to heavily penalize moving beyond a 35km radius from the center point. By doing this, CEs move only within proximity of the enemy to 'attack' if the enemies move into the CEs' sector. To generate the path, the new position is randomly selected between the four nodes minimizing the distance measure, and one random node within the 5% smallest measures. The probability for selecting one of the four smallest measures is set twice as high as the one node belonging to the 5% smallest. Furthermore, the 4 previously visited nodes are excluded by blacklisting those positions, except the current node, such that a CE can remain static.

Capacity

Each CE is equipped for 48 hours of operations, but the required amount of each class of supplies depends on the specific role. Every CE corresponds with a platoon. Based on each role for this unit size, an average supply capacity is estimated, displayed in Table 5-1. These estimates are a rough representation based on classified data of actual supply consumption.

Generating consumption

Supply consumption is modeled to be dependent on the distance to the closest CE. It is assumed that maximum consumption is reached within a 2km proximity, and decreases linearly up to 30km. Except for fuel, the average consumption is defined to be the capacity divided over 48 hours, which is consumed at 16km distance to the closest enemy. For fuel, the known consumption per kilometer is multiplied by the traveled distance since the past period. Not all classes vary with the same degree on this proximity. The dependency is estimated to be 0% in class 1 (water and food), 35% in class 2 (equipment), 25% in class 3 (fuel), 35% in class 4 (materials and spare parts) and 90% in class 5 (munitions). For example, equipment is then consumed at a rate of 65% of the mean value at a range of 30km or larger, and at 135% at 2km or less.

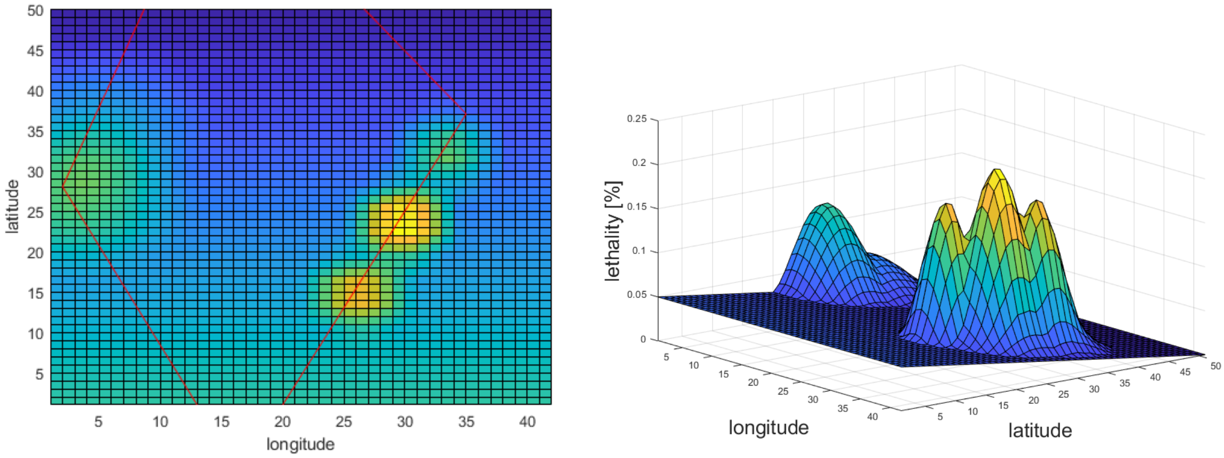

(a) The AOO with axes of advance

(b) Cubic interpolated lethality.

Figure 5-2: A mapping of the risks associated with the advancing enemies.

Thresholds & weights

It is estimated that a CE requires sufficient inventory to reach and maintain its most intense operation in the coming 6 hours for an additional 6 hours. For example, when its maximum consumption is planned in 4 hours it needs the planned consumption up until that point and sufficient inventory to maintain that maximum consumption for 6 hours, even if it has planned to retreat. As weights, uniformly random variables are taken in the interval $[0.2, 1]$, as it is difficult to provide any sensible, general estimate of the tactical weights represented.

Generating risks

In this scenario, the west axis consists of lighter, and the east one of heavier enemies. The advance takes 72 hours, though each enemy starts with a varying delay of at most four hours. In Table 5-2 the different types and their properties are displayed. Each type has an inner range, outer range, and lethality. The lethality describes the probability that a vehicle is destroyed over a period of 1 hour, at the exact position of the enemy. At the edge of the inner range, the lethality is reduced to 20% of its original value, and at the outer range it is completely diminished. Then, a grid is sampled spaced at one kilometer, by cubic interpolation of the lethality at the center, inner and outer ranges. Furthermore, a linear risk decreasing from 20% at the enemy side of the front, to 0% at the own side of the front, describes the long-range artillery and air power. The overall risk is set as the maximum of all risks at that location. Using this, a mapping can be created as displayed in Figure 5-2a.

Table 5-2: The parameters for each type of enemy

Type	Inner range	Outer range	Lethality	Axis	Amount
Reconnaissance	10km	20km	25%	west	1
Light mechanized	8km	15km	50%	west	1
Heavy mechanized	5km	10km	70%	east	2
Armored	6km	12km	90%	east	1

Table 5-3: Transit risk example for three sectors that need to be crossed.

sector	A	B	C	Actual transit risk
Risk / hour	10%	20%	5%	$1 - 0.9 \cdot 0.8 \cdot 0.95 = 31.6\%$
Risk / 30 min	5%	10%	2.5%	$1 - 0.95 \cdot 0.9 \cdot 0.975 = 16.6\%$
Risk / 20 min	3.3%	6.7%	1.7%	$1 - 0.967 \cdot 0.933 \cdot 0.975 = 11.3\%$

Defining transit risks

To retrieve the static risks described in Chapter 3, the interpolated risk can simply be selected at the node's location. The transit risks, however, are more complicated. As the risks describe the chance of being destroyed each hour, it is dependent on the travel time of each vehicle. If sectors are crossed in less time, the risk is lower for each crossed sector, which has a cumulative effect. To prevent calculating an intractable amount of transit risks for each separate vehicle, the transit risks are normalized to single period transits. This is illustrated in Table 5-3. A normalization factor $n = \frac{d_{ij}}{3}$ is defined, with d_{ij} being the distance between node i and $j \in N$, and 3km the fixed sector size in which a strike can occur. The transit risks for varying speeds can then quickly be approximated using the normalized transit risks. For example, a travel time of 1.5 hours would mean $1 - (1 - 0.113)^{1.5} = 16.4\% \approx 16.6\%$, and for a travel time of 3 hours $1 - (1 - 0.113)^3 = 30.2\% \approx 31.6\%$.

Available vehicles

Supply can be performed by different Robotic Autonomous Systems (RAS), or in the classical approach by manned convoys. In Table 5-4 the properties of three types of vehicles and the classical convoy are listed. For each type, the speed, fuel and supply capacity, range, and profile are defined. The vehicle's profile describes a constant decreasing both the static and transit risks linearly. The larger the profile, the bigger the chance of being destroyed. As described in Chapter 3, the inventory capacity is measured per supply class. It is assumed that each vehicle carries a capacity equal to the mean capacity of all CEs per class.

Table 5-4: Properties of the available vehicles, consisting of two typical types of RAS, a truck that can be both manned or autonomous, and the classical convoy.

Vehicle	Type	Speed	Fuel cap	Profile	Range	Supply cap
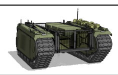	UGV	25km/h	30L	0.25	300km	1.200kg
	Cargo UAV	100km/h	8L	0.5	100km	250kg
	Truck	50km/h	200L	1	500km	4.000kg
	Convoy	50km/h	600L	1	500km	12.000kg

Chapter 6

Results

The main question from the problem statement (see Section 1-4) is formulated as:

Can an efficient solution for DARE be developed such that dispersed combat operations can be effectively supported while remaining scalable in computation time and sufficiently robust against unreliable communication?

As is highlighted in Section 2-3-3, it is possible to achieve both scalability and robustness against communication failures by adopting a specialized coordinative framework for partitioning the problem space into smaller (overlapping) subsets. To be able to implement such a framework, the solution method needs to allow coordination by adopting the Multi-Robot System (MRS) paradigm for Distributed Constraint Optimization Problem (DCOP), such that the solution process is distributed over all vehicles, as is explained in Section 4-1-1.

Proving that such a coordinative framework contributes to robustness and scalability, is out of scope for this thesis (see the related thesis by Korthals Altes (2021)). However, it still remains to be shown that the distributed solution approach for the model describing Dispersed Autonomous Resupply (DARE) is actually an efficient solution method that returns qualitative solutions, meaning that dispersed operations are effectively supported. To measure the effectiveness, the indicators from Section 1-3-1 are used as performance criteria. The quantitative measure for each indicator is displayed below:

Indicator	**Qualitative measure**
Responsiveness	Maximum penalty over time
Sustainability	Total penalty accumulated
Flexibility	Invariance against sudden changes in Combat Element (CE) plans
Survivability	Invariance against increasing risks of destruction
Efficiency	Total vehicles and supplies lost

The main questions that need to be answered, can then be formulated as:

- Does a RAS network consisting of Unmanned Ground Vehicles (UGVs) and Unmanned Areal Vehicles (UAVs) perform the DARE concept better than a classical convoy?
- Do the proposed local and cooperative algorithms contribute significantly to flexibility and survivability?
- To which degree does distributed cooperation influence the result, and how is this divided between MGM and MGM-2?
- What is the effect of more rounds of cooperation, and how many rounds could reasonably be achieved given the computation time?
- To which degree are the results dependent on this specific scenario instance?

To analyse the effectiveness of executing the DARE concept using a Robotic Autonomous Systems (RAS) network, the experiments are simulated using both a RAS network and a classical convoy. As is highlighted in Table 5-4, a convoy is modelled as one vehicle but corresponds with the parameters of three trucks combined. The standard parameter setup for the experiments is the following:

Parameter	**Value**
Number of runs	8
Simulation time	24 hours
Period length	15 minutes
Time horizon	6 hours
Optimization rounds	3
Plan update frequency	$f_p = 1/2\text{h}$
CEs	1x reconnaissance, 1x engineer, and 2x infantry
Inventory	$I_{kc,0} = 25\% \cdot b_{kc} \quad \forall c \in \mathcal{C},\ k \in K$
RAS network	4 UGVs, 2 UAVs & 1 truck

The plan update frequency dictates after how many periods a new plans with risks, CE positions, consumption, thresholds and weights is provided to the vehicles. This means that new data is generated following the steps described in Chapter 5. Furthermore, as initially the CEs are uniformly distributed over the Area Of Operations (AOO), and the distance from any Forward Operating Base (FOB) to the AOO is roughly equal, the vehicles are initialized at a random FOB in all simulations.

All experiments are performed on an HP Elitedesk with a 3.8GHz AMD A12-9800E processor.

6-1 Effectiveness of using RAS

Responsiveness & sustainability

In the first experiment, the scores for each of the performance criteria are collected for the standard parameter setup. As the fleet size and mix influence the performance, it is decided to also simulate three separate trucks. The capacity of these trucks is equal to the convoy and, just like the RAS network, it can perform cooperation and separate deliveries.

In Table 6-1 the results are displayed for the RAS network, a single convoy, and three separate trucks. As can be clearly seen, the RAS approach significantly outperforms the convoy approach on all performance criteria. Not only is a smaller portion of the fleet lost with fewer supplies, but also the responsiveness and sustainability are better achieved by using RAS. Of course, this effect is dependent on when the convoy is destroyed. Possibly, the convoy outperforms the RAS setup when still alive.

Table 6-1: Average results of 8 simulation runs for the standard parameters.

		Penalties		Losses	
Setup	Capacity	Max	Total	Vehicles	Supplies
No supply	-	1.301	305.926	-	-
1 convoy	12.000kg	824	63.843	50%	3.436kg
3 trucks	12.000kg	630	26.235	46%	3.148kg
7 RAS	9.300kg	579	22.768	21%	1.334kg

In Figure 6-1 a boxplot shows the spreading of the results of all simulations. Here, it is visible that RAS performs consistently better than the convoy approach. However, it does not seem to perform significantly better than the case with three separate trucks. This indicates that it is not the capacity, but the flexibility and redundancy that dictate performance. In general, the RAS network removes 55% of the maximum penalty and 93% of the total penalty compared to the case without any resupply, which is a substantial increase in responsiveness and especially sustainability.

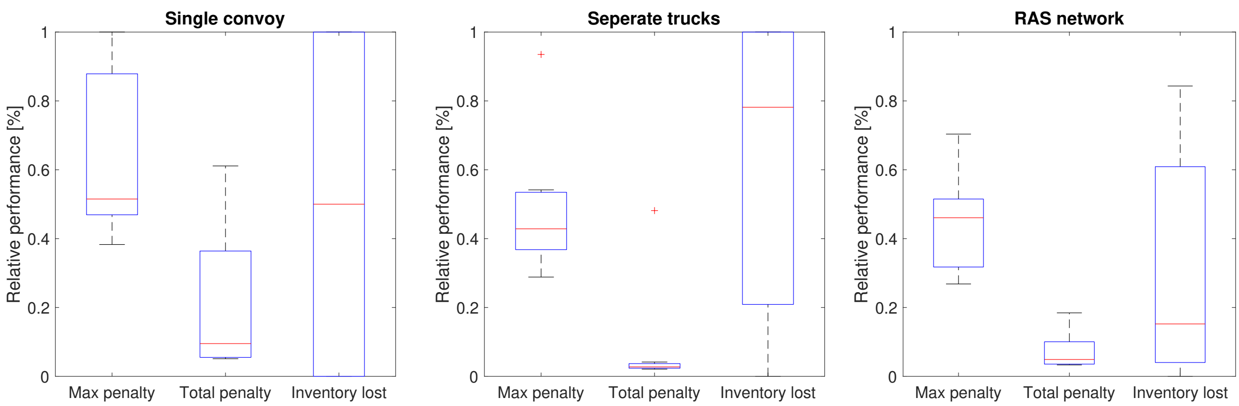

Figure 6-1: The relative performance for each setup. As can be seen, RAS performs more consistent and significantly better over all performance criteria compared to convoys.

Survivability

To test for survivability, some degree of invariance against increasing risks of destruction should be visible. Therefore, in the second experiment, the risks are varied. To be specific, the lethality of each enemy, as described in Section 5-2, is reduced by a linear factor:

Parameter	Value		
$x \cdot lethality$	$x = 0.01$	0.5	1.0

The latter case, with $x = 1.0$, corresponds with the results already obtained in the previous experiment. The results are displayed in Table 6-2. On average, RAS keeps performing better. However, the difference in survivability decreases rapidly, and RAS even shows bigger losses with very low risks. This makes sense, as a larger amount of vehicles results in a larger probability of at least one vehicle being destroyed.

Table 6-2: Average results of 8 simulation runs using varying risks.

Case	Setup	Penalties Max	Penalties Total	Losses Vehicles	Losses Supplies
$x = 0.01$	1 convoy	566	18.236	0%	0kg
	7 RAS	581	13.694	2%	24kg
$x = 0.5$	1 convoy	1.096	33.136	22%	1.433kg
	7 RAS	572	20.199	17%	598kg
$x = 1.0$	1 convoy	824	63.843	50%	3.436kg
	7 RAS	579	22.768	21%	1.334kg

What is interesting from these results, is that responsiveness and sustainability seem unaffected for the RAS setup by increasing risk, though supply and system losses have increased. This indicates good survivability. The effect is depicted in Figure 6-2. What is also apparent, is that with low risk the setups perform comparable, supporting the idea that the main benefit of using a RAS network is enhanced survivability.

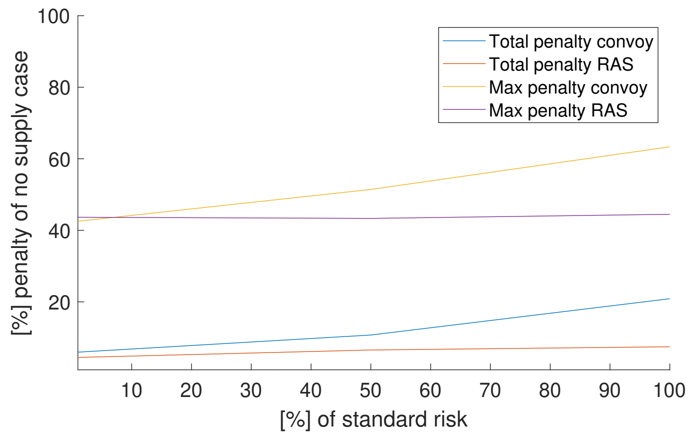

Figure 6-2: A comparison of the responsiveness and sustainability against varying risks.

Flexibility

Along with survivability, the performance in flexibility is measured by invariance against sudden changes in the provided information. Therefore, also an experiment is performed with a smaller update frequency. Lastly, both a low update frequency and a low risk are combined.

Parameter	**Value**		
f_p	$1/24h$	$1/6h$	$1/2h$

As $f_p = 1/2h$ is the standard parameter, this result corresponds with the first experiment. The results are displayed in Table 6-3. It is visible that both setups profit from a lower update frequency, which can be expected. The best performance can be observed when both the update frequency and the risk are reduced. When there is lower flexibility required, performance is comparable, but when the frequency goes up, the differences diverge quickly. This could be explained by the simple numbers with which RAS is deployed. If plans are changed, it is more likely that vehicles find themselves in the right spot at the right time.

Table 6-3: Average results of 8 simulation runs using varying plan update frequencies.

		Penalties		Losses	
Case	Setup	Max	Total	Vehicles	Supplies
$f_p = \frac{1}{24},\ x = 0.01$	1 convoy	560	8.925	0%	0kg
	7 RAS	583	7.567	0%	0kg
$f_p = \frac{1}{24}$	1 convoy	576	9.932	30%	2.640kg
	7 RAS	552	8.768	16%	659kg
$f_p = \frac{1}{6}$	1 convoy	636	15.524	44%	2.773kg
	7 RAS	607	20.251	16%	1.394kg
$f_p = \frac{1}{2}$	1 convoy	824	63.843	50%	3.436kg
	7 RAS	579	22.768	21%	1.334kg

In Figure 6-3 the relative performance is displayed. Again, the RAS setup shows almost constant performance indicating good flexibility, which is a clear difference from the convoy.

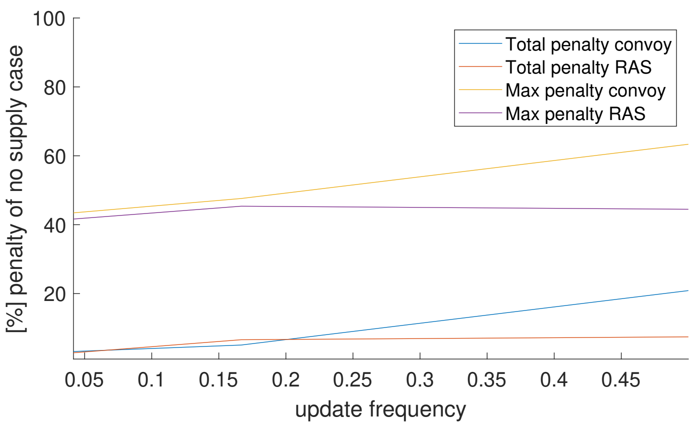

Figure 6-3: A comparison of the invariance against varying update frequencies.

6-2 Effect of distributed cooperation

Types of cooperation

In the next experiment, the effect of distributed cooperation is analysed. The RAS network is tested without any cooperation, with only Maximum Gain Messaging (MGM) and with both MGM and MGM-2, along with the convoy setup for comparison. In Figure 6-4 the results are displayed in a boxplot. In total, the effect of cooperation on the maximum and total penalty seems only marginal. However, it is interesting to see that survivability is actually the lowest with cooperation. From the previous experiments in Section 6-1 it became clear that survivability negatively influences the performance criteria, meaning that with cooperation the performance was marginally better *despite* having lower survivability. This means that the positive effect might be larger than can be observed directly. A possible explanation for the lower survivability is that without cooperation vehicles only select deliveries discounted by their expected survivability. During cooperation, this discount is not used. Hence, if 'safe' deliveries are already performed, riskier deliveries might be performed instead. However, including the risk discounts could potentially also have an adverse effect, as more vehicles need to schedule a risky delivery to prevent a penalty from occurring, such that more vehicles are drawn towards risks.

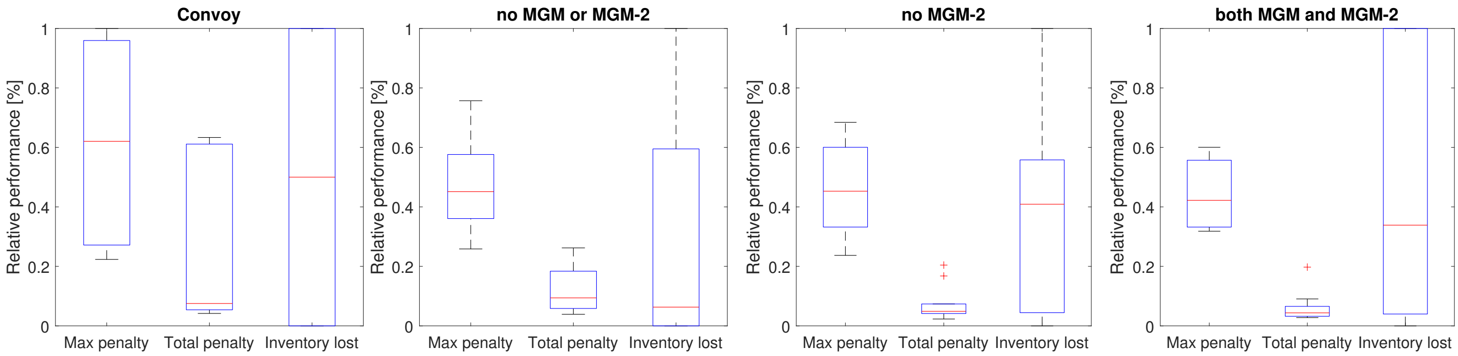

Figure 6-4: Relative performance for the cooperation setups.

In Table 6-4 the average results are displayed for all cooperation setups. Overall, the combination of MGM and MGM-2 seems to perform best, though not to a large extent.

Table 6-4: Average results of 10 simulation runs for different cooperation setups.

		Penalties		Losses		Transfers	
Case	Run time	Max	Total	Vehicles	Supplies	Fuel	Supplies
No resupply	-	1310	308.156	-	-	-	-
Convoy	254s	789	73.027	50%	3.326kg	-	-
No MGM or MGM-2	936s	607	36.215	16%	1.369kg	-	-
MGM	4340s	585	22.760	26%	1.609kg	66kg	1719kg
MGM & MGM-2	6173s	568	18.979	21%	1.192kg	54kg	980kg

Apart from performance, also the total amount of transferred resources is displayed. It is interesting to see that the volume of transfers has gone done down by using MGM-2. This suggests that MGM-2 mainly successfully removes constraints, instead of proposing them. However, in Figure 6-5, it becomes clear that this is not the case. A possible explanation might be that vehicles with a large fuel capacity accept many constraint proposals to refuel UAVs, such that there remain fewer possibilities for larger supply transfers.

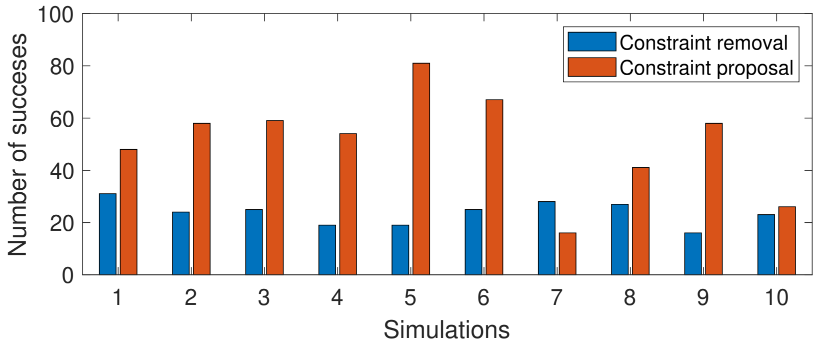

Figure 6-5: Number of successes per MGM-2 type for each simulation run.

Duration of cooperation

Furthermore, simulations are performed with varying numbers of cooperation rounds. Within each round, both MGM and MGM-2 are performed sequentially. Each cooperation type is therefore performed the same number of rounds.

Parameter	**Value**		
Rounds	1	3	6

In Table 6-5 the results are displayed, including the run times. Surprisingly, the performance of 6 rounds is significantly worse than with 3 rounds. In the setup with 6 rounds, however, the volume of supply and fuel transfers is significantly higher. A logical explanation would be that the more rounds are added, the more probable that vehicles will execute dangerous transfers. This is possible, as the probability for survival is not included when selected transfers.

Table 6-5: Average results of 8 simulation runs using a different number of rounds for cooperation.

	Run Time		Penalties		Losses		Transfers	
Case	Total	Per period	Max	Total	Vehicles	Supplies	Fuel	Supplies
No resupply	-	-	1317	311.136	-	-	-	-
Convoy	236s	2.5s	932	93.211	57%	3.832kg	-	-
1 round	876s	9.1s	741	51.626	20%	2.498kg	-	-
3 rounds	6308s	66.7s	578	21.138	16%	750kg	81kg	2.375kg
6 rounds	13199s	137.5s	725	36.808	22%	1.322kg	119kg	4.809kg

On average, performing 3 rounds of both MGM and MGM-2 combined takes roughly a minute. This means that, with the given period length of 15 minutes, in the real-time case up to 45 rounds could be executed with this algorithm. Though this is promising, the results do not indicate a positive effect of performing those extra rounds.

In Figure 6-6 also a boxplot is included to display the spread of the results. Though on average the case with 6 rounds performs only marginally better than a single round, the losses seem more consistent.

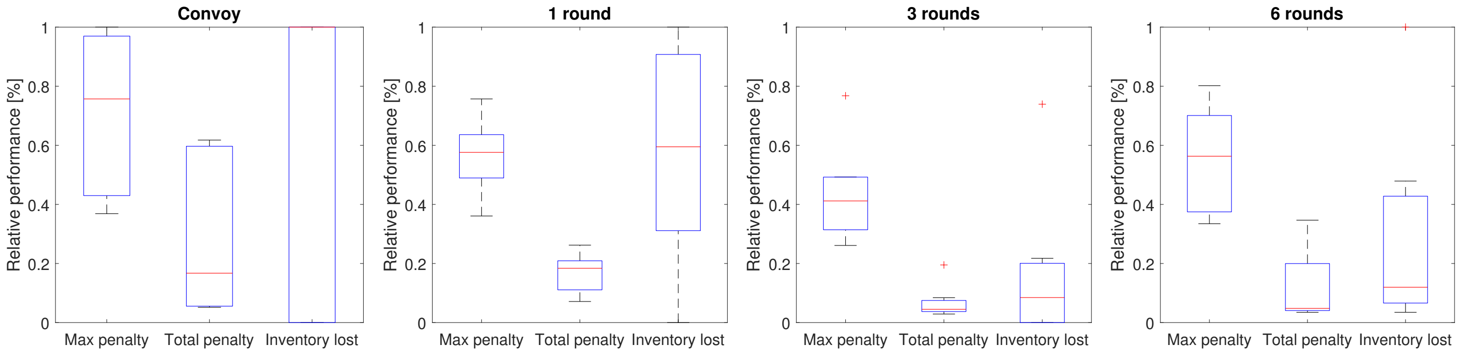

Figure 6-6: Relative performance for setups with a different number of rounds.

6-3 Dependency on simulation instance

The question remains whether the observed performance increase of RAS over convoys is not too dependent on the chosen scenario instance. This cannot be excluded completely and is possibly even to be expected. In the end, this model is specifically designed for a concept of operation, and not for all military missions thinkable. Still, to verify the results, the scenario data is relaxed to compare performance. To achieve this, the random perturbations of the enemy paths are enlarged. As described in Chapter 5, the actual position of the enemy $\{l'_y, l'_x\}$ is randomized by perturbing the path of the axis $\{l_y, l_x\}$ with a normal distribution, using $\{\sigma_y, \sigma_x\} = \{1.2, 0.2\}$. The perturbations are enlarged using the following factors:

Parameter	**Value**		
$x \cdot \{\sigma_x, \sigma_y\}$	$x = 1$	3	5

Upon these locations, the positions and thus consumption of the own CEs are heavily dependent. The effect of the enlarged perturbations is made visible in Figure 6-7. In these plots, all generated positions of all plans over a complete simulation horizon are plotted. This means that, over a time span of 24 hours, with an update frequency of 1/2h, a total of 12 (overlapping) plans are plotted. Obviously, the 5 times larger perturbations do not show a very specific scenario, as opposed to the original perturbations, where the axes of advance are still clearly visible.

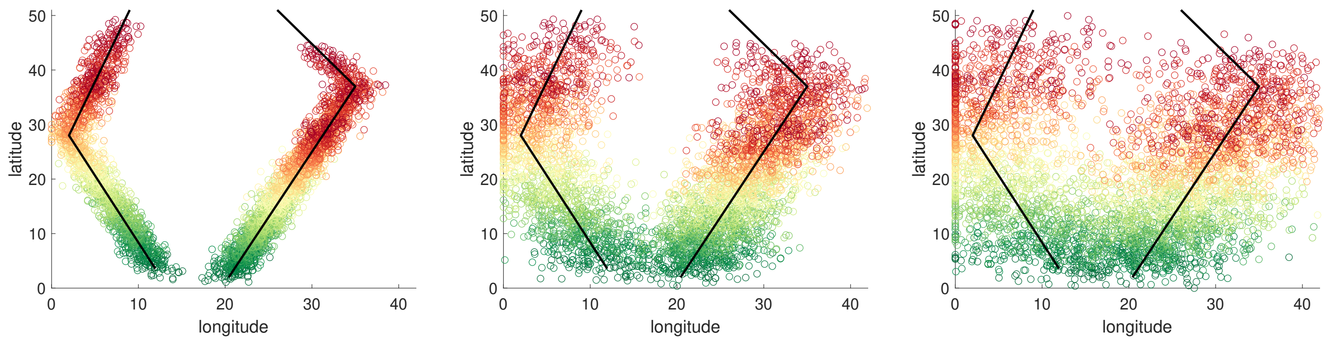

(a) 1x perturbation **(b)** 3x perturbation **(c)** 5x perturbation

Figure 6-7: Different scales of perturbations for the enemy paths, with each dot being an enemy positions, and each color representing a different update.

The results of running simulations based on these input data are displayed in Table 6-6. As can be seen, performance for both the convoy and RAS remain comparable between the standard case and the 3x larger perturbations. The 5x perturbations result in significantly higher penalties and less survivability. Overall, RAS still outperforms the convoy, suggesting that the obtained results are not completely dependent on this scenario instance.

Table 6-6: Average results of 8 simulation runs using varying perturbations of enemy paths.

		Penalties		Losses	
Case	Setup	Max	Total	Vehicles	Supplies
$x = 1$	1 convoy	824	63.843	50%	3.436kg
	7 RAS	579	22.768	21%	1.334kg
$x = 3$	1 convoy	801	79.379	50%	3088kg
	7 RAS	601	27.111	19%	937kg
$x = 5$	1 convoy	808	81.257	64%	4.382kg
	7 RAS	688	43.245	29%	1.547kg

Chapter 7

Conclusions

The current global balance of power requires the army to maintain itself in high-intensity combat against a peer opponent. To do so, it seeks to apply innovative concepts like dispersed operations. However, this new concept poses complex challenges for logistics, for which currently no viable solutions exist. To this end, the new concept Dispersed Autonomous Resupply (DARE) is developed, based on Robotic Autonomous Systems (RAS). For the first time, an integrated approach is proposed to deal with lacking information, a dynamic environment, and the risk of enemy attacks. Because of this novelty, a critical analysis is required.

Overall, a RAS based setup for DARE appears to outperform the classical convoy approach. It has been shown that under low risk or low dynamic circumstances the performance is equal between the two setups, but that RAS shows much larger flexibility and survivability, with near-complete invariance against increasing risks and plan dynamics. In reality a convoy would not be used to perform DARE, as already has been stated in Section 1-2. But because improving the classical approach is found to be hardly viable, the shown flexibility and survivability indicate that RAS is an interesting alternative.

The effect of adding cooperation is positive, though it is not found to be consistent and in many cases only marginal. However, as is shown in Section 6-2, the cooperative approach showed marginally better performance despite lower average survivability. This suggests that there are unknown implicit effects of cooperation, which is supported by the finding that more cooperation resulted in worse performance. An effect could for example be that extra cooperation leads to vehicles taking more risk to perform transfers, as no risk is included when scheduling those. On the other hand, the same cooperation might also prevent multiple vehicles from trying to perform a highly beneficial, but risky delivery. Though this effect is positive, it can also lead to insufficient redundancy in deliveries, such that when a vehicle is destroyed, a large penalty might occur as no other vehicle can quickly back it up. Overall, there seems to be room for improvement in cooperation.

The generality of the results is explored by testing multiple perturbations of enemy paths, as analysed in Section 6-3. Though the results are fairly consistent, concluding that the results are not dependent on the scenario is a far stretch. Apart from the enemy paths, there are many more factors influencing performance. A very important factor is for instance the vehicles' risk profile, for which only initial estimates are used. However, the risk of being destroyed can be hardly assumed constant. This is very dependent on the available armament and sensors of the enemy units, and can as such heavily influence survivability. To mitigate this, pessimistic estimates are used, but more research with accurate data and dynamic risk profiles would be needed to verify the consistency of the results. The same issue occurs with the resulting penalties and the corresponding fleet mix. In the analysis, only the relative performance is compared. It is not assessed in an absolute sense if the performed deliveries would be sufficient to proceed with the planned operations. As a consequence, the required fleet size and mix might be much different from the used example, of 4 Unmanned Ground Vehicles (UGVs), 2 Unmanned Areal Vehicles (UAVs), and one truck for 4 platoons. It seems reasonable to expect that a different ratio between for example UGVs and UAVs would generate different results. Again, to verify the consistency, accurate data would need to be used along with a qualitative military assessment.

Future work can for example be aimed at including dynamic risk profiles during cooperation, such that vehicles with a low profile are favored when having to perform risky deliveries. Furthermore, the concept of synchronized rounds of cooperation could be relaxed. This would allow algorithms to deal better with realistic communication constraints. Next to cooperation, also the local algorithms could be improved. An integrated approach to scheduling and path planning leads to fewer restraints during scheduling, and thus better delivery plans. On top of that, the use of exact methods might be feasible for subsets of the problem, such as the rebalancing of the vehicles. Furthermore, one of the major issues remaining is the complexity of the labeling algorithm for path planning. Possibly, the algorithm itself can be improved by including relaxed dominance rules, better approximations of the expected survival probability, or by adding heuristical approaches to limit the set of nodes to a set of expected relevant nodes. For both the local and cooperative algorithms, no effort is taken yet to escape local minima. Possibly, some concept of simulated annealing or multi-start could aid performance. In addition, in both cases the performance might be improved by reusing information from past optimizations. Currently, each time the information is updated, the whole optimization is restarted, losing any good but unaffected partial solutions. On top of the algorithmic suggestions, also the model could be improved. For instance, currently vehicles are required to wait until the end of the period to perform cooperation, which could be relaxed to increase flexibility. Also the added possibility for UGVs to carry UAVs might produce better results, as currently the range of the UAV severely limits their use.

To truly be able to answer the problem statement, coordination would also have to be included. This would require to implement a coordinative method, along with the corresponding analysis of the effect of specific neighborhoods on performance. Furthermore, a wider range of experiments and scenarios is required. An interesting venue would be to compare the effect of RAS with convoys in a more classical combat scenario. Also the benefit in different types of (peacekeeping) missions is interesting. Furthermore, terrain limitations have not been implemented yet, which is often a severe problem in realistic combat scenarios.

Concluding, this thesis can be considered an exploratory study on the integration of a new optimization model with RAS, for which the results already show consistent and promising flexibility and survivability. Though there is an area for improvement, this work can serve both as a groundwork and supporting argument for further development of a RAS-based concept for tactical logistics.

References

13e Lichte Brigade. (2018). *Operationeel concept lichte brigade.* Koninlijke Landmacht.

Abbas, H., Shaheen, S., & Amin, M. (2015). Organization of multi-agent systems: An overview. *International Journal of Intelligent Information Systems*, *4*(3), 46-57.

Ackerman, E. (2014, February 12). *Lockheed's robotic trucks pass real-world military convoy test.* Retrieved 2020-02-17, from https://spectrum.ieee.org/automaton/robotics/military-robots/lockheeds-robotic-trucks-pass-real-world-military-convoy-test

Ackerman, E. (2017, February 1). *Swarms of disposable drones will make critical deliveries and then vanish.* Retrieved 2020-02-17, from https://spectrum.ieee.org/automaton/robotics/drones/otherlab-apsara-aerial-delivery-system

Ackerman, E. (2019, October 22). *Zipline deploys medical delivery drones with u.s. military.* Retrieved 2020-02-17, from https://spectrum.ieee.org/tech-talk/robotics/drones/zipline-successfully-deploys-medical-delivery-drones-with-us-military

Agatz, N., Bouman, P., & Schmidt, M. (2018). Optimization approaches for the traveling salesman problem with drone. *Transportation Science*, *52*(4), 965–981.

Alberts, D. S. (2007). *Agility, focus, and convergence: The future of command and control* (Tech. Rep.). Office of the Assistant Secretary of Defense for Networks and Information.

Baker, S., & Shi, P. (2002). Formulation of a tactical logistics decision analysis problem using an optimal control approach. *ANZIAM Journal*, *44*, 82–113.

Barahona, F., Chowdhary, P., Ettl, M., Huang, P., Kimbrel, T., Ladanyi, L., … Swirszcz, G. M. (2007). Inventory allocation and transportation scheduling for logistics of network-centric military operations. *IBM Journal of Research and Development*, *51*(3.4), 391-407.

Berbeglia, G., Cordeau, J.-F., & Laporte, G. (2010). Dynamic pickup and delivery problems. *European journal of operational research*, *202*(1), 8–15.

Bertazzi, L., Savelsbergh, M., & Speranza, M. G. (2008). Inventory routing. In *The vehicle routing problem: latest advances and new challenges* (pp. 49–72). Springer.

Bettinelli, A., Ceselli, A., & Righini, G. (2011). A branch-and-cut-and-price algorithm for the multi-depot heterogeneous vehicle routing problem with time windows. *Transportation*

Research Part C: Emerging Technologies, *19*(5), 723–740.

Bouros, P., Sacharidis, D., Dalamagas, T., & Sellis, T. (2011). Dynamic pickup and delivery with transfers. In *Advances in spatial and temporal databases* (pp. 112–129). Springer.

Chauhan, D., Unnikrishnan, A., & Figliozzi, M. (2019). Maximum coverage capacitated facility location problem with range constrained drones. *Transportation Research Part C: Emerging Technologies*, *99*, 1–18.

Church, R., & ReVelle, C. (1974). The maximal covering location problem. In *Papers of the regional science association* (Vol. 32, pp. 101–118). Springer-Verlag.

Coltin, B., & Veloso, M. (2014, 09). Online pickup and delivery planning with transfers for mobile robots. *Proceedings - IEEE International Conference on Robotics and Automation*, 5786-5791.

Dell'Amico, M., Montemanni, R., & Novellani, S. (2019). Models and algorithms for the flying sidekick traveling salesman problem. *arXiv preprint:1910.02559*.

Dorling, K., Heinrichs, J., Messier, G. G., & Magierowski, S. (2016). Vehicle routing problems for drone delivery. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, *47*(1), 70–85.

Drexl, M. (2012). Synchronization in vehicle routing—a survey of vrps with multiple synchronization constraints. *Transportation Science*, *46*(3), 297–316.

Drexl, M. (2013). Applications of the vehicle routing problem with trailers and transshipments. *European Journal of Operational Research*, *227*(2), 275 - 283.

Edwards, S. J. A. (2005). *Swarming and the future of warfare* (PhD Dissertation). Pardee RAND Graduate School, Santa Monica, CA.

Farahani, R. Z., Abedian, M., & Sharahi, S. (2009). Dynamic facility location problem. In *Facility location* (pp. 347–372). Springer.

Feillet, D., Dejax, P., Gendreau, M., & Gueguen, C. (2004). An exact algorithm for the elementary shortest path problem with resource constraints: Application to some vehicle routing problems. *Networks: An International Journal*, *44*(3), 216–229.

Ferrandez, S. M., Harbison, T., Weber, T., Sturges, R., & Rich, R. (2016). Optimization of a truck-drone in tandem delivery network using k-means and genetic algorithm. *Journal of Industrial Engineering and Management (JIEM)*, *9*(2), 374–388.

Fragkos, I., Cordeau, J.-F., & Jans, R. (2017). *The multi-period multi-commodity network design problem.* CIRRELT, Centre interuniversitaire de recherche sur les réseaux d'entreprise.

Gallasch, G. E., Lilith, N., Billington, J., Zhang, L., Bender, A., & Francis, B. (2008). Modelling defence logistics networks. *International Journal on Software Tools for Technology Transfer*, *10*(1), 75–93.

Gue, K. R. (2003). A dynamic distribution model for combat logistics. *Computers & Operations Research*, *30*(3), 367 - 381.

Ham, A. M. (2018). Integrated scheduling of m-truck, m-drone, and m-depot constrained by time-window, drop-pickup, and m-visit using constraint programming. *Transportation Research Part C: Emerging Technologies*, *91*, 1–14.

Ivanova, K., Gallasch, G. E., & Jordans, J. (2016). *Automated and autonomous systems for combat service support: scoping study and technology prioritisation.* Defence Science and Technology Group Edinburgh SA Australia.

Jeong, H. Y., Song, B. D., & Lee, S. (2019). Truck-drone hybrid delivery routing: Payload-energy dependency and no-fly zones. *International Journal of Production Economics*, *214*, 220–233.

Kablau, R. (2002). Het fysieke distributieconcept van de koninklijke landmacht. *Militaire Spectator*, *171*, 540 - 551.

Karber, P. A. (2015). Lessons learned from the russo-ukrainian war. *The Potomac Foundation.*

Keirin, J. (2016). Logistics support for unmanned systems. In *Operations research for unmanned systems* (p. 233-253). John Wiley & Sons, Ltd.

Klibi, W., Lasalle, F., Martel, A., & Ichoua, S. (2010). The stochastic multiperiod location transportation problem. *Transportation Science*, *44*(2), 221-237.

Koninklijke Landmacht. (2017). *Studie slimme en robuuste logistiek.* Ministerie van Defensie.

Korthals Altes, M. A. (2021). *A framework for intent-based command and control of robotic autonomous systems for persistent reconnaissance* (Unpublished master's thesis). Delft Center for Systems & Control.

Maheswaran, R. T., Pearce, J. P., & Tambe, M. (2004). Distributed algorithms for dcop: A graphical-game-based approach. In *Isca pdcs* (pp. 432–439).

Marufuzzaman, M., Nur, F., Bednar, A., & Cowan, M. (2020). Enhancing benders decomposition algorithm to solve a combat logistics problem. *OR Spectrum.*

Masson, R., Lehuédé, F., & Péton, O. (2013). An adaptive large neighborhood search for the pickup and delivery problem with transfers. *Transportation Science*, *47*(3), 344–355.

Murray, C. C., & Chu, A. G. (2015). The flying sidekick traveling salesman problem: Optimization of drone-assisted parcel delivery. *Transportation Research Part C: Emerging Technologies*, *54*, 86 - 109.

Otto, A., Agatz, N., Campbell, J., Golden, B., & Pesch, E. (2018). Optimization approaches for civil applications of unmanned aerial vehicles (uavs) or aerial drones: A survey. *Networks*, *72*(4), 411–458.

Parragh, S. N., Doerner, K. F., & Hartl, R. F. (2008). A survey on pickup and delivery models part ii: Transportation between pickup and delivery locations. *Journal für Betriebswirtschaft*, *58*(2), 81–117.

Prodhon, C., & Prins, C. (2014). A survey of recent research on location-routing problems. *European Journal of Operational Research*, *238*(1), 1–17.

Pugliese, L. D. P., & Guerriero, F. (2013). Solution approaches for the elementary shortest path problem.

Rabta, B., Wankmüller, C., & Reiner, G. (2018). A drone fleet model for last-mile distribution in disaster relief operations. *International Journal of Disaster Risk Reduction*, *28*, 107–112.

Ren, J., Zheng, X.-l., & Tan, Y.-j. (2013). Multi-product supply planning for combat units in battlefield environment. In *The 19th international conference on industrial engineering and engineering management* (pp. 445–453). Springer.

Roldán, R. F., Basagoiti, R., & Coelho, L. C. (2017). A survey on the inventory-routing problem with stochastic lead times and demands. *Journal of Applied Logic*, *24*, 15 - 24.

Sebbah, S., Ghanmi, A., & Boukhtouta, A. (2011). Modeling and simulation of military tactical logistics distribution. In *Proceedings of the 2011 winter simulation conference (wsc)* (p. 2502-2513).

Sebbah, S., Ghanmi, A., & Boukhtouta, A. (2013). A column-and-cut generation algorithm for planning of canadian armed forces tactical logistics distribution. *Computers & Operations Research*, *40*(12), 3069 - 3079.

Shavarani, S. M., Nejad, M. G., Rismanchian, F., & Izbirak, G. (2018). Application of hierarchical facility location problem for optimization of a drone delivery system: a case study of amazon prime air in the city of san francisco. *The International Journal of Advanced Manufacturing Technology*, *95*(9-12), 3141–3153.

Skubch, H. (2013). *Modelling and controlling of behaviour for autonomous mobile robots.* Springer.

Tambe, M. (1997). Towards flexible teamwork. *Journal of artificial intelligence research*, *7*, 83–124.

Thornton, S., & Gallasch, G. E. (2018). Swarming logistics for tactical last-mile delivery. In *International conference on science and innovation for land power.*

Toyoglu, H., Karasan, O. E., & Kara, B. Y. (2011). Distribution network design on the battlefield. *Naval Research Logistics (NRL)*, *58*(3), 188–209.

UK Ministry of Defence. (2017, June 29). *Competition document: autonomous last mile resupply.* Retrieved 2020-02-04, from https://www.gov.uk/government/publications/accelerator-competition-autonomous-last-mile-supply

US Marine Corps. (2016). *Logistic operations.* Department of the Navy.

US Marine Corps. (2018). *Tactical-level logistics.* Department of the Navy.

Walker, A. (2018, April 9). *Autonomous last mile resupply: Titan robot put through its paces.* Retrieved 2020-02-17, from https://www.qinetiq.com/blogs/2018/04/autonomous-last-mile-resupply

Wang, X., Poikonen, S., & Golden, B. (2017). The vehicle routing problem with drones: Several worst-case results. *Optimization Letters*, *11*(4), 679–697.

Glossary

List of Acronyms

ADS	Ammunition Distribution System
AOO	Area Of Operations
APC	Armoured Personnel Carrier
CPN	Coloured Petri Net
C2	Command & Control
CE	Combat Element
CONOPS	Concept of Operations
CSS	Combat Service Support
DARE	Dispersed Autonomous Resupply
DCOP	Distributed Constraint Optimization Problem
DFLP	Dynamic Facility Location Problem
DoD	Department of Defence
ESPPRC	Elementary Shortest Path Problem with Resource Constraints
FOB	Forward Operating Base
FSTSP	Flying Sidekick Traveling Salesman Problem
IFV	Infantry Fighting Vehicle
IRP	Inventory Routing Problem
ISR	Intelligence, Surveillance & Reconnaissance
LAV	Light Armoured Vehicle
LCE	Logistics Combat Element
LMD	Last-Mile Delivery
LRP	Location Routing Problem
MAS	Multi-Agent System
MCNDP	Multi-Commodity Network Design Problem

MCFLP	Maximum Coverage Facility Location Problem
MGM	Maximum Gain Messaging
MOB	Main Operating Base
MRS	Multi-Robot System
NCO	Network-Centric Operation
PDPT	Pickup-and-Delivery Problem with Transfers
PDPOT	Pickup-and-Delivery Problem with Online Transfers
PDSTSP	Parallel Drone Scheduling Traveling Salesman problem
POD	Point of Debarkation
POE	Point of Embarkation
RAS	Robotic Autonomous Systems
RFS	Reconnaissance Fire System
SA	Situational Awareness
TSPD	Traveling Salesman Problem with Drone
UAV	Unmanned Areal Vehicle
UGV	Unmanned Ground Vehicle
VRP	Vehicle Routing Problem
VRPD	Vehicle Routing Problem with Drones
VRPPD	Vehicle Routing Problem with Pickup and Delivery
VRPTT	Vehicle Routing Problem with Trailers and Transshipment