Titre: Title:	Agent-Based Systems and Dynamic Multi-Agent Scheduling for Fleet Management in Underground Mines: Towards Mining 4.0
Auteur: Author:	Giuseppe Basilico
Date:	2020
Type:	Mémoire ou thèse / Dissertation or Thesis
Référence: Citation:	Basilico, G. (2020). Agent-Based Systems and Dynamic Multi-Agent Scheduling for Fleet Management in Underground Mines: Towards Mining 4.0 [Mémoire de maîtrise, Polytechnique Montréal]. PolyPublie. https://publications.polymtl.ca/9153/

URL de PolyPublie: PolyPublie URL:	https://publications.polymtl.ca/9153/
Directeurs de recherche: Advisors:	Jean-Marc Frayret, Michel Gamache, & Diane Riopel
Programme: Program:	Maîtrise recherche en génie industriel

POLYTECHNIQUE MONTRÉAL
affiliée à l'Université de Montréal

Agent-Based Systems and Dynamic Multi-Agent Scheduling for Fleet Management in Underground Mines: Towards Mining 4.0

GIUSEPPE BASILICO
Département de mathématiques et de génie industriel

Mémoire présenté en vue de l'obtention du diplôme de *Maîtrise ès sciences appliquées*
Génie industriel

Août 2021

POLYTECHNIQUE MONTRÉAL
affiliée à l'Université de Montréal

Ce mémoire intitulé :

Agent-Based Systems and Dynamic Multi-Agent Scheduling for Fleet Management in Underground Mines: Towards Mining 4.0

présenté par **Giuseppe BASILICO**
en vue de l'obtention du diplôme de *Maîtrise ès sciences appliquées*
a été dûment accepté par le jury d'examen constitué de :

Camélia DADOUCHI, présidente
Jean-Marc FRAYRET, membre et directeur de recherche
Michel GAMACHE, membre et codirecteur de recherche
Diane RIOPEL, membre et codirectrice de recherche
Michel RIOUX, membre

DEDICATION

To my parents Antonio and Concetta

ACKNOWLEDGEMENTS

I would like to express my deep gratitude to my supervisors Prof. Jean-Marc Frayret, Prof. Michel Gamache and Prof. Sergio Terzi, for giving me the opportunity to embark on this mind-blowing 2-year experience and for providing me with ongoing inspiration and support throughout the whole research project. I would like to thank Politecnico di Milano for allowing me to participate in the double degree programme in Canada, a journey that made me grow both as an engineer and as a person. I wish to extend my thanks to all the friends I met in Canada, for making this experience one of the happiest times of my life. Last but not least, special thanks go to my parents Antonio and Concetta and my brother Davide, for their love, endless support and encouragement.

RÉSUMÉ

Cette thèse présente un système multi-agents et une stratégie d'ordonnancement dynamique multi-agents pour la gestion de la flotte d'équipements dans les mines souterraines. Les activités de gestion de la flotte visées, plus spécifiquement, sont le dispatching, le routage et la gestion du trafic des véhicules miniers, qui concernent respectivement: l'attribution de la destination suivante à un véhicule qui vient de terminer une tâche; le choix de l'itinéraire à suivre pour atteindre la destination visée; la coordination du trafic dans le réseau de transport souterrain, qui se compose de segments de tunnel bidirectionnels à une seule voie. Afin de reproduire le comportement et d'évaluer les performances du système de gestion de la flotte à base d'agents proposé, un modèle de simulation à base d'agents d'une mine d'or souterraine canadienne est conçu en utilisant AnyLogic. Les résultats de l'étude de simulation indiquent que, en adoptant la solution proposée dans ce mémoire, le système minier serait en mesure d'atteindre des volumes de production plus élevés, une plus grande capacité d'expansion du réseau souterrain, et une meilleure réactivité à la demande des zones de travail souterraines qu'actuellement. Ceci est synonyme de meilleure efficacité du système logistique et productif. De plus, on observe une amélioration de la performance du goulot d'étranglement du système, une meilleure utilisation des ressources et une diminution du nombre de kilomètres parcourus par tonne de minerai extrait. Ceci est synonyme de meilleure efficience du système logistique et productif. Enfin, grâce au système de gestion de flotte multi-agents, la mine bénéficie d'une agilité et d'une auto-organisation accrues. La solution proposée s'intègre dans le paradigme Industrie 4.0 et constitue les premiers pas vers la “mine intelligente”. Dans cette optique, une des contributions de ce mémoire est de détecter les améliorations ponctuelles des indicateurs clés de performance de la mine, et ainsi de quantifier les retours sur l'investissement dans l'Industrie 4.0.

ABSTRACT

This thesis presents an agent-based system and a dynamic multi-agent scheduling strategy for fleet management in underground mines. The fleet management activities addressed are dispatching, routing and traffic management of mining vehicles, which deal respectively with: the assignment of the next destination to a vehicle that has just completed a task; the choice of the route to be followed to reach the selected destination; the coordination of vehicle traffic in the underground transportation network, made up of one-lane bi-directional tunnel segments. To evaluate the proposed solution, an agent-based simulation model of a Canadian underground gold mine is built with AnyLogic. By implementing the agent-based system incorporating the dynamic multi-agent scheduling strategy in the digital replica of the mine, it is in fact possible to reproduce its behaviour and evaluate its performance. In the light of the results of the simulation study, the designed agent-based fleet management system proves to outperform the one currently employed by the mine under investigation. The main key performance indicators, related on the one hand to the effectiveness and on the other to the efficiency of the logistics and production system, are in fact found to be improved. Adopting the proposed solution would ultimately mean embracing the Industry 4.0 paradigm and taking the first steps towards the "smart mine". Against this background, the main achievement of this study is to capture precise improvements in mining key performance indicators, and thus quantify the returns on investment in Industry 4.0.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF SYMBOLS AND ACRONYMS

ABM	Agent-based modelling
ABS	Agent-based simulation
LHD	Load-haul-dump vehicle
AGV	Automated guided vehicle
FMS	Fleet management system
KPI	Key performance indicator
DynCNET	Dynamic contract net (protocol)
CNET	Contract net (protocol)
FIFO	First in first out

LIST OF APPENDICES

CHAPTER 1 CONTEXT AND PROBLEM DEFINITION

1.1 Introduction

Underground mines are industrial plants whose objective is to extract minerals from underground ore deposits, transport the extracted ore to the surface, and process the raw materials into finished products for the upstream tiers of the mining supply chain.

This research focuses on the extraction and transportation processes, with an emphasis on the management of the respective production and logistics systems. A key aspect, in particular, will be tackled: the management of the fleet.

The term *fleet management* usually refers to all management activities related to a company's vehicle fleet. These include activities of financial nature (e.g. choice to purchase or lease vehicles), activities related to vehicle maintaining (e.g. maintenance), activities related to vehicle running (e.g. refuelling), and activities of logistical and operational nature. The research topic is therefore quite multifaceted. This thesis, like all the other works in the underground mining literature dealing with the same subject, will address the latter class of management activities.

Delving into details and following the convention used in [6], managing the fleet of an underground mine means dealing with three types of decision:

1. Dispatching decision
   Assign the next destination to a resource that has just finished its current task.

2. Routing decision
   Define the route to be followed to reach the chosen destination point.

3. Scheduling decision
   Define the vehicle movement so as not to have traffic conflicts with other vehicles moving in the mine. This decision is necessary because of the nature of the underground transportation network, which consists of one-lane bi-directional road segments.

Depending on the management policy adopted by the specific mining company, the frequency and the complexity of the dispatching problem may vary. If the firm opts for a rolling approach, as is most often the case due to the high operational uncertainty inherent in underground mining, the decision will have to be taken whenever a vehicle finishes its current task. The resources schedule, therefore, will be built dynamically over time. If, on the contrary, the company decides to solve a single major initial problem, in which tasks are assigned to resources for a specific time frame (usually not going beyond the working day, again due to operational uncertainty), the decision will have to be made at the same time for each vehicle, considering the entire planning horizon. Each mining resource will thus be provided with a schedule in which the activities to be carried out are specified for the entire time window under consideration. Routing and scheduling decisions, on the other hand, are usually delegated to vehicle drivers. They are thereby taken dynamically during the execution of logistics and production processes.

The dispatching, routing and scheduling issues forge what is known in the underground mining literature as the *fleet management problem*. This research project aims to propose an innovative solution to this problem, employing agent-based systems and dynamic multi-agent scheduling, and embracing the Industry 4.0 paradigm.

1.2 Underground mining processes

In order to better outline the profile of the fleet management problem in an underground mine, it is appropriate to analyse the underground mining processes. The following is a general overview.

Two main processes take place in an underground mine: development mining and production mining. Development mining means excavating in order to gain access to the orebody, while production mining means excavating and extracting ore from the orebody which has been accessed and made available by development. There are several different methods of production mining, but they will not be further investigated for the purposes of this research. As far as development mining is concerned, on the other hand, the individual activities forming part of the process will be analysed hereafter. The goal is to create underground tunnels by excavating valuable ore and non-valuable material (waste). In order to do this, development cycles are executed iteratively, progressively increasing the length of the created tunnel. Figure 1.1 depicts the development cycle, detailing which activities need to be performed and which resources (i.e. vehicles) are required for each activity. It is worth noting that operations *washing* and *shotcreting* may not be part of the process in some underground mines. The individual operations take place at the end of the underground tunnel, called face. In a nutshell, each cycle aims to separate material from the current face, thus lengthening the tunnel, and to prepare the tunnel for the next cycle.

First, holes are drilled by drilling vehicles and charged with explosives by emulsion trucks. Then the blast takes place, and the ensuing toxic fumes are ventilated. Subsequently the

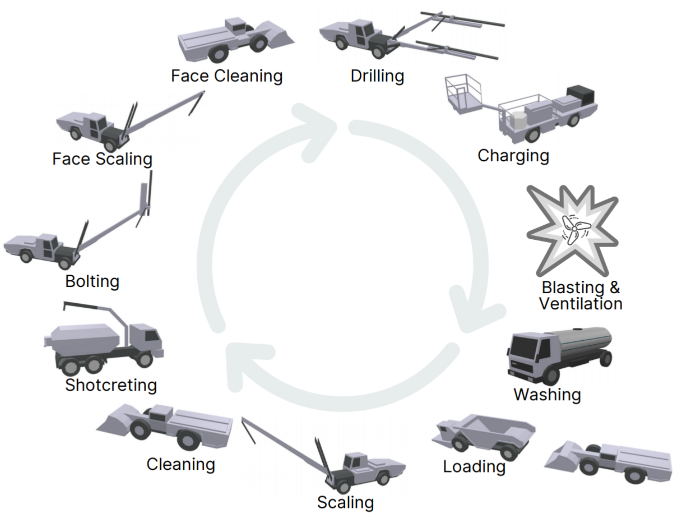

Figure 1.1 The development cycle: process, activities and resources

face is sprayed with water in order to reduce the number of airborne particles. The rock that was separated as a result of the explosion is removed from the face by the so-called load-haul-dump (LHD) vehicles; they load the bulk into trucks, which can thus begin their journey towards the Earth's surface, where the freight will be unloaded and will become the raw material of the mining industrial plant. In the next step of the cycle a scale rig scales off loose rock still attached to the interior of the tunnel. The rock scaled away is then removed. In some underground mines a shotcreter then uses concrete to reinforce the walls and the roof, preventing the tunnel from collapsing. Further reinforcement is achieved in the following operation, where the walls and ceiling are bolted to the surrounding rock mass. Finally, the last two finishing activities scale the face and remove the rock scaled away. The development cycle has thus come to an end: the underground tunnel has become longer and is ready for the next cycle.

For safety reasons, all workers must be outside the mine for the explosion to take place. The blast is therefore scheduled in time windows called *blast windows* (underground mines usually work 12-hour shifts, and the time slots dedicated to blasting are those elapsing between two successive shifts). When the blast window is approached, all mine operations are paused, the workers leave the mine, the explosion takes place, and finally the next shift workers return to the underground work areas, resuming the interrupted activities. The only exception to business interruption concerns shotcreting, which cannot be paused since the concrete would cure. This operation needs therefore to be scheduled in-between blast windows.

1.3 Fleet management problem profile

Having presented on the one hand an introduction to the problem addressed by this research, and on the other an overview of the main underground mining processes, we now have at our disposal all the theoretical tools necessary to accurately outline the fleet management problem profile in underground mines. In the following, in this respect, the research topic will be analysed in terms of decision variables, constraints and objectives.

As already hinted in Section 1.1, for each vehicle, there are three decision variables in the problem: destination (dispatching decision); road (routing decision); coordination mechanism with other travelling vehicles (scheduling decision). Looking at the decision-making problem from another perspective, it is immediate to grasp the fact that it is about deciding which resource does what, where and when. The fleet management problem, therefore, boils down to what in Operations Management is known as the *scheduling problem.*

As far as the constraints of the problem are concerned, it is worth pointing out that the logistics and production system to be managed features numerous logistics, process and safety constraints. First of all, on account of the high infrastructure costs of tunnelling, the haulage network consists exclusively of one-lane bi-directional road segments, with the presence of occasional passing bays where vehicles can temporarily stop to allow oncoming vehicles to pass. This severely restricts vehicle movements, complicating the problem to be solved. The traffic coordination of vehicles in the transportation network is therefore the first constraint to be met. With regard to working areas, on the other hand, underground sites require to be visited sequentially by different types of vehicles, each performing a specific operation (Figure 1.1): the need for coordination of the activities carried out at each extraction point, and consequently the need for coordination of the tasks assigned to each resource, are therefore of paramount importance. Another constraint concerns the interruption of operations due to safety requirements during blast windows. This must be taken into account when scheduling activities and assigning tasks to resources, as previously explained (the shotcreting operation, for example, cannot be interrupted and must therefore be scheduled in-between blast windows). The underground mining literature, to conclude, is awash with mentions of the high operational uncertainty, with constant references to unexpected events, to the high randomness of operations times, and so on and so forth. Further details will be provided during the literature review chapter, but for now it is important to note that this is an additional pitfall to proper fleet management.

Finally, as far as the objective function of the fleet management problem is concerned, the mining company must aim to maximise the efficiency and effectiveness of its logistics and production system. According to the author, priority should be given to effectiveness. In the mining sector, in fact, demand rules the global mineral supply and demand equation [7]: especially in the case of mines producing gold, such as the one investigated in this thesis, demand usually exceeds supply, and firms are sure to sell what they produce. In such a market environment, effectiveness should take precedence over efficiency. Thereupon the first objective a mine must pursue when managing its production and logistics system is the maximisation of production volumes. But, of course, striving for one objective does not mean completely neglecting the other. From the costs side, in fact, the importance of a proper fleet management is underlined by the fact that transportation and extraction costs account for about 40% of the total operating costs of an underground mine [8].

Having outlined the decision variables, the constraints, and the objective function of the fleet management problem in underground mines, the reader should now have in mind an accurate picture of the problem addressed by this research.

To summarise in a nutshell what has been discussed so far, one can say that the fleet management problem in underground mining, which can be traced back to the scheduling problem in Operations Management, consists of managing and coordinating a multitude of heterogeneous vehicles, which must obey numerous process and safety constraints, moving in a transportation network with very binding characteristics, having as background an environment marked by a high operational uncertainty. Such a problem is a major logistical challenge, and this thesis will aim to face it by leveraging agent-based systems and dynamic multi-agent scheduling, and taking the first steps towards the "smart mine" envisaged by the Industry 4.0 paradigm.

The remaining of this thesis proceeds as follows. Chapter 2 presents the literature review. Scientific contributions, research objectives and a theoretical framework are discussed in Chapter 3, followed by the methodology in Chapter 4. After this, there is a presentation of a real underground mine (industrial partner) and its digital replica created by the author with AnyLogic. The agent-based fleet management system proposed by this thesis is then presented in Chapter 6. Finally, the last three chapters present the results of the study, a discussion and a conclusion.

CHAPTER 2 LITERATURE REVIEW

2.1 Introduction

This introductory section will explain the purpose of the literature review and the boundaries of analysis. Last but not least, gaps in the literature will be highlighted.

The objective of this literature review is to list and analyse existing fleet management systems (FMSs) for underground mines, with particular emphasis given to the related dispatching, routing, and scheduling problems.

Pursuing a broader perspective and considering the entire management of the logistics and production system of an underground mine, FMSs are integrated downstream of the whole planning process, after the strategic, tactical and operational planning. Although the underground mining literature is full of publications proposing tools for these three upstream phases of the planning process, these papers will not be reviewed as they are outside the focus of this research. The reader, however, should be aware that some of the works concerning short-term operational planning deal with the dispatching decision. In other words, they propose tools capable of generating short-term plans in which not only certain activities to be carried out during the upcoming shift(s) are scheduled, but in which such activities are also already assigned to the fleet vehicles. Since the dispatching problem is one of the fundamental pillars of this thesis, these works, although outside the field of FMSs and related to operational planning, will be considered in the following discussion. Table 2.1 below outlines the most important articles belonging to this class of “outsiders”, highlighting for each of them the most relevant features for the purposes of this discussion.

Table 2.1 A collection of short-term operational planning tools whose output already includes the dispatching decision

		[9]	[10]	[11]	[12]	[13]
Short-term plan generated with the proposed tool	Planning horizon	About a couple of days (the most imminent upcoming shifts)	Upcoming working shift	Upcoming working shift	2 months	Not explicitly defined (can range from monthly to daily plans)
	Time unit (for planning horizons broader than 1 shift)	Hour			Shift	Not defined
	Are the activities to perform in each underground site scheduled over the planning horizon?	Yes	Yes	Yes	Yes	Yes
	Are such activities assigned to resources? In other words, is the dispatching decision settled?	Yes	Yes	Yes	Yes	Yes
Nature of the activities considered		Activities of the entire development cycle and transportation	Activities of the entire development cycle and transportation	Activities of the entire development cycle and transportation	Only transportation	Only activities of the development cycle
Nature of the resources dispatched		LHDs[1] and all other vehicles used throughout all the development cycle	Vehicles for both transportation and the development cycle	Vehicles for both transportation and the development cycle	LHDs - Trucks	Vehicles for carrying out the development cycle
Are operational, legal, safety aspects/constraints considered?		Yes	Yes	Yes	No	Yes
Approach to handle stochasticity/unexpected events, if mentioned		Event-based rescheduling[2]			Event-based rescheduling	Event-based rescheduling
Information related to the hierarchical planning approach (strategic, tactical and operational planning)	Input data to feed the tool (from the upstream decision-making process)	Extraction plan determining the tonnage to be extracted. Number of development cycles to be performed in each underground site during the panning horizon	Extraction plan derived from a monthly plan	Extraction plan	Monthly production target for each draw point	
	Downstream decision-making process fed with the tool output	Routing			(optionally) Fleet Management System / Dispatch System	

LEGEND

[1] *LHDs: load-haul-dump vehicles*

[2] *Event-based rescheduling : rescheduling at unexpected events (possible thanks to quick solution times) in order to incorporate the latest known information in a rolling horizon approach; re-running the tool whenever underground operating conditions change with respect to the forecast*

As pointed out in the table, the reviewed articles propose tools capable of producing operational plans in which mining operations are on the one hand scheduled for the next shift(s), and, on the other hand, allocated to resources. When employing such approaches, the dispatching decision is not taken by the FMS, but by the short-term planning tool, one step ahead in the planning hierarchy. On an operational level, following a more pragmatic perspective, each vehicle would already have at the beginning of the upcoming shift a list of destinations to visit during the next shift(s). Such strategies have been negatively criticised by, among many others, [6], [14] and [15]. Making plans over a long period of time such as the working shift, in fact, is something inappropriate in the context of an underground mine, due to the high level of uncertainty (unexpected events, highly stochastic nature of operations, etc.); variable conditions in the mine require, the other way round, real-time decision making, in order not to degrade the performance of the haulage fleet and the productivity of the mining system. The reader must be aware, moreover, that the vision of this research is to propose a FMS capable of handling the stochastic nature of the underground mine and of responding efficiently and effectively to unexpected events. Assigning at the beginning of the shift activities and destinations for the entire shift(s) is a strategy that, promoting a rigid system, moves in the opposite direction to the research objectives. As a conclusion, the articles presented in Table 2.1, both because promoting inappropriate planning philosophies for an underground mine [6,14,15], and because conflicting with the vision of this thesis, will not be further investigated, even if able to solve the dispatching problem.

Focusing on the real core of this research, i.e. FMSs and the related dispatching, routing, and scheduling problems in underground mining, the reader should acknowledge that, at least from an academic point of view, little progress has been made in recent years, and considering the current state of the art there is a scarcity of models and algorithms. [6, 14, 16–20] are witnesses to this poor literature. The scarcity of academic solutions for the fleet management in underground mines can be defined high, especially when the benchmark is the number of existing solutions for the same fleet management problem in the different mining context of open-pit mines. It should be noted, in fact, that the FMSs literature for open-pit mines is full of academic contributions (the interested reader may refer to [21], a review that gives a well-rounded representation of the current state of the art of this topic, considering both academic contributions and commercially implemented FMSs). The gap in the progress and innovation status of FMSs for the two types of mines can be explained by the facilitated application of traffic monitoring systems (which ease the use of FMSs) in open-pit mines with respect to underground mines, and, more importantly, by the lower complexity of the fleet management problem in open-pit mining: in an underground mine, in fact, the constrained haulage network composed of one-lane bidirectional road segments require three decisions to be made (dispatching, routing, scheduling), contrary to the single dispatching decision to be taken for an open-pit mine. It should also be noted that an academic solution or a commercial fleet management software intended for open-pit mines cannot be applied in underground mines, due to the diversity of the two management problems; evidence of the non-transferability of solutions from one type of mine to the other can also be found in [14–17]. The void in the literature on underground mining cannot therefore be filled with the thriving literature on open-pit mining.

In a nutshell, there are many gaps in the current literature of FMSs in underground mining, and this opens up opportunities for academic researches like this to advance the state of the art of this topic.

In the rest of this chapter, always targeting the underground mining sector, the main works dealing with FMSs and the related dispatching, routing, and scheduling problems will be reviewed.

2.2 State of the art - Academic contributions

The aim of this section is to present the reader with a critical survey of the main academic publications concerning FMSs and the related dispatching, routing, and scheduling decisions in underground mines.

A real-time fleet management system that can be used to take dispatching, routing, and scheduling decisions whenever a vehicle, upon reaching the destination and completing the related task, asks for a new assignment, is the *shortest path algorithm* proposed by [6].

This tool was inspired by the procedures presented in [22], employable in the manufacturing field to decide which is the route that allows an AGV (automated guided vehicle) to reach a known destination in the shortest possible time, without having traffic conflicts with other AGVs in a bidirectional flow path network. [22]'s approach, dealing with the routing and scheduling aspects of the problem within the same solution process, was adapted by [6] in order to be transferable to the underground mining world. New aspects of the fleet management problem were introduced, in order to have a solution process capable of ensure proper vehicle orientation at destination (bucket in front), reduce deadlock occurrences, and include also the dispatching decision along with the routing and scheduling ones.

The following is an overview of the *shortest path algorithm.* The approach consists in applying the Dijkstra's algorithm to a time window graph, built as follows. Once the transportation network of the underground mine has been virtually divided into segments, at each end-point of each segment will correspond *m*-*n* nodes of the time window graph, where *m* is the number of instants of time (discretized) from the moment of the request to the end of the planning horizon, and *n* is the number of times the end-point is already occupied by other vehicles in the same time span. Two nodes of the time window graph are connected by an arc only if geometric, time, conflict (considering both catching-up and head-on collisions), and side-stepping tests are verified (see [22] for more information). In order to apply the Dijkstra's algorithm, weights must be placed on the arcs of the time window graph, and this is done according to the dispatching criteria the decision maker wants to follow. Two approaches are presented by [6]: in the former, weights represent the travel time in the road segments; in the latter, weights are computed to consider the deviation of the production levels from the target levels of the optimal production plan. Whenever a vehicle that has finished its assignment asks for a new destination (to be chosen from a known set of possible destinations representing loading/dumping points), the time window graph is dynamically constructed by applying the Dijkstra's algorithm and considering all the schedules assigned in the past to other vehicles. As output of the *shortest path algorithm*, the vehicle will be dispatched to a new loading/dumping point: in the case of the first approach, it will be assigned to the destination that can be reached in the shortest possible time; in the case of the second approach, it will be assigned to the destination corresponding to the minimum deviation from the optimal production plan. In both cases, in addition to the destination, the vehicle will be notified of the road to be travelled with the respective times, so as to be compliant with the traffic of the other vehicles already scheduled. Launching Dijkstra's algorithm on the time window graph means, therefore, taking dispatching, routing and scheduling decisions simultaneously.

This fleet management tool, despite being able to solve the three problems at the same time, presents two weaknesses, as recognized by the authors of the approach in their paper.

1. First of all, we are talking about a myopic approach: when a vehicle asks for a new task, the three management decisions are taken solely and exclusively for the requesting vehicle, without taking into account the fact that in a more or less far future there will be new vehicles requesting a new task. The approach, therefore, is selfish, and results in performances of the transportation system worse than those achievable with a more collaborative approach, the aim of which would not be to find an optimal local plan for the requesting vehicle, but a good global solution for the whole system.
2. A second key drawback of the algorithm is the consideration of all input data as deterministic values. Among the inputs of the technique there are: the schedule of vehicles assigned in the past; the travel time of the requesting vehicle in each segment of the network; the loading/unloading time of the demanding vehicle. As is easy to guess, considering determin-

istic travel, loading and unloading times means feeding the algorithm with estimates of the actual time. But in reality these times are stochastic variables, many times characterized by a high degree of variability due to unexpected events such as equipment breakdowns, unplanned stops in a drift, etc. This means that, in the real word, the algorithm, which claims to calculate conflict-free travel schedules, by considering deterministic times will not achieve its purpose, giving rise to conflicting schedules embodied in the transportation system in the form of traffic problems. In order for the algorithm to no longer guarantee the creation of conflict free schedules, in particular, it is sufficient that at least one of the actual travel, loading and unloading times is greater than its estimate used as input to the algorithm[1]. If, the other way around, the actual times were lower than the estimates, it is true that there would be no traffic problems, but the algorithm would produce local solutions no longer optimal. The problem just exposed is further aggravated by the fact that the schedule of a vehicle, output of the algorithm, is the input of the decision making process of all the vehicles that will be dispatched in the future: if the vehicle does not respect its assigned schedule due to actual times different from those expected, as soon as a deviation is detected, in order not to have sub-optimal solutions or conflicts, the algorithm should be relaunched for all the vehicles scheduled after that vehicle. In conclusion, if at least one of the travel, loading and unloading times deviated from those expected for a vehicle, the real time fleet management system would not only make bad decisions for the vehicle in question, but would compromise all decisions made for vehicles scheduled after that vehicle.

Moreover, in addition to the two problems abovementioned, the proposed real-time FMS has the drawback of not taking into account the entire development cycle. In fact, the *shortest path algorithm* was created to manage only and exclusively the transportation fleet (more specifically the load-haul-dump vehicles), and therefore does not consider all the other activities and vehicles of the development cycle (Figure 1.1). The problem thus considered is far removed from the real problem faced by underground mines, and the solution found with the tool, being related to a simplified problem, would be difficult to implement in a real underground mine.

Thanks to the tests performed in [14], we know that the [6]'s real-time FMS improves the KPIs of the transportation system compared to the techniques usually used in underground

---

[1] A possible solution to the "non-guarantee" of conflict-free schedules was proposed by [22], and consists in introducing a safety factor to increase the expected time, thus reducing the risk of the traffic problems mentioned above (this risk would be nullified if the algorithm was re-launched as soon as the vehicle did not respect the schedule comprehensive of safety factors [22]). This, however, would not lead to an overall improvement of the FMS performance, as the reduction of traffic problems would be in trade-off with the goodness of the dispatching, routing and scheduling solutions. This shortening, moreover, was also criticized by [6].

mines. However, as explained above, the tool has the weaknesses of being a myopic approach and not considering both the stochastic nature of mining operations and the whole development cycle. The FMS proposed in this research will try, among other things, to solve these three criticalities.

[16] is undoubtedly one of the pioneers of the study of dispatching, routing and scheduling problems in underground mines. By proposing a strategy for the management of a fleet of remote-controlled/automatic load-haul-dump vehicles (RAL), [16] was the first to address specific traffic management problems, such as, for example, vehicle motion in bidirectional lane-segments. Before 1991, most of the work on RAL systems in underground mines was hardware oriented, and no one had ever specifically discussed the three fleet management problems of dispatching, routing and scheduling. These three topics, meanwhile, had been investigated in the manufacturing industry, thanks to the progressive introduction of AGVs in company logistics. None of the associated studies, however, analysed the three problems for a system composed of bidirectional paths, which is the reference network to which the layout of an underground mine can be traced. [16], therefore, can be considered the forerunner of the exploration of the themes treated in this research.

The approach proposed by [16] focuses mainly, if not exclusively, on the scheduling part of the problem, not presenting particular and refined solutions for the dispatching and routing aspects. It could be used for the fleet management in an underground mine as follows. Once a vehicle is available for a new assignment, all possible destinations are considered (the choice between them, i.e. the dispatching decision, will be made only at the end of the procedure). For each destination, the Dijkstra's algorithm is applied considering the time as a criterion, thus identifying the path to reach the destination in the shortest possible time. A first critique to this approach can be made. The path so found, in fact, may not be in reality the shortest one in term of time, because at this stage only the vehicle in question is considered, without worrying about previously scheduled vehicles moving in the mine. This is one of the problems that occurs when routing and scheduling are not executed simultaneously in the same resolution process. In [16]'s approach, in fact, the scheduling part of the problem is considered only after a route for each destination has been identified, and works as follow. Considering deterministic travel, loading and unloading times, a provisional schedule is generated. At this point the compatibility of the provisional schedule with the schedules of all vehicles assigned in the past is evaluated, and three algorithms are applied to resolve any possible conflict: an algorithm that controls whether a vehicle should slow down due to another slower vehicle travelling in the same direction; an algorithm handling bidirectional conflicts; an algorithm handling passages across traffic zones. By applying these algorithms, the provisional schedule of the vehicle under consideration is modified, so that a conflict-free route is found to arrive at the destination. If a solution to a conflict is not found, the possibility to reschedule the conflicting vehicle already assigned in the past is evaluated. This procedure, potentially, can call into question all the schedules found in the past. If a conflict is unsolvable even considering the rescheduling of vehicles allocated in the past, the destination concerned is discarded. As a result of all these steps, conflict-free roads will be obtained for each possible destination. The destination that can be reached in the shortest possible time is chosen.

The solution presented by [16] presents the same problems as [6]: it is a myopic approach, it considers deterministic times, and it does not take into account the whole development cycle. It seems to have a lower computational effort, and therefore a shorter response time, than [6]'s *shortest-path algorhitm*, due to the fact that only the scheduling part of the problem is considered in depth. The fact of "neglecting" the routing and the dispatching, by the way, seems to reduce the quality of the solution found.

Relying on dynamic programming concepts, [17] proposed an enumeration algorithm capable of, fixed the dispatching decisions, optimally solving routing and scheduling problems in underground mines.

The paper can be seen as a resumption and extension of the concepts in the article [23], capable of considering, in addition, the displacement mode of vehicles: for this reason, between the two publications, only [17] will be discussed throughout this review.

Assuming that when a vehicle becomes free at a loading or dumping point its next destination is already known, the algorithm is able, considering the vehicles that are moving in that moment toward given destinations, to optimally compute the conflict-free path able to lead each vehicle to its destination. For optimality we refer to the minimization of the time elapsed from the moment of demand by the vehicle to that in which the final state, in which each vehicle has reached its destination, is reached.

The approach is based on the concept of state, defined as a set of vectors able to define the orientation and the exact position of each vehicle in the haulage network. A state, more in detail, defines for each vehicle: the segment in which it is positioned (for each segment representing a road of the haulage network, a forward and a return virtual segment are defined); the time remaining to reach the end of the segment; a binary variable that defines the orientation of the vehicle.

When a vehicle requests the launch of the routing and scheduling algorithm, the starting state is defined, and all possible states in which at least one vehicle changes segment with respect to the initial state are generated from it. At this point, states are eliminated and/or updated thanks to the execution of some tests able to: identify intersection crossing conflicts; identify catching-up and head-on conflicts; modify states in order to make them admissible; take into account the orientation of the vehicle and eliminate states where the orientation is not proper. In addition, a filtering procedure is performed to eliminate states dominated by others or not respecting certain criteria, in order to reduce the computational effort of the algorithm (for more detailed information the reader can consult [17]). Among the remaining son states of the initial state, the one closest to the final state is selected (closeness is defined by the sum of the times related to the shortest paths of each vehicle, neglecting potential conflicts), and its child states are generated, tested and filtered, following the same procedure described above. Among the states that have survived the testing and filtering process and whose children have not yet been produced, the one closest to the final state is selected, and its son states are generated, tested and filtered. This last procedure is iterated, resulting in a progressive expansion of the enumeration tree. The computation stops as soon as the tree englobes the final state, in which the vehicle that asked which path to follow and the related timetable in order to reach a given destination, as well as all the other vehicles that were moving at that time, arrive at their destinations. The sequence of states in the tree going from the initial state to the final one provides each vehicle with the conflict-free route to follow to reach its destination.

Like all the approaches already presented, [17] has the limitations of being a myopic technique, considering deterministic times, and being limited to the transportation aspect instead of managing the entire development cycle. However, the negative consequences resulting from the first two weaknesses are partly mitigated by the greater dynamism of this approach compared to the previous ones: at each decision, in fact, the algorithm re-optimizes the routes of all vehicles in the system. The attenuation effect could be even greater if the enumeration algorithm was combined with a dispatching decision-making process with a high level of dynamism.

[24] proposed a tool that could be used to take scheduling decisions in a ramp of an underground mine. The authors formulated a mixed integer programming model able to optimize vehicle schedules in order to maximize the throughput of the transportation system, expressible by the number of trips performed during a shift (one vehicle trip corresponds to leaving the mine surface, travelling to the loading point, loading, returning to the surface, unloading).

The schedules are created in such a way that laden trucks ascending the ramp do not have to stop in the passing bays, and traffic conflicts are avoided thanks to temporary vehicle stops (in the bays, on the surface before descending the ramp, at the loading point before ascending the ramp). Only head-on scheduling conflicts are treated. Conflicts in intersections are not addressed, since the model does not consider the entire transportation network but only the ramp (featuring no intersections). Catching-up collisions are not considered either, since all vehicles have the same speed.

The optimization tool has many simplifications and limitations which restrict its sphere of application. First of all, it is able to schedule vehicle movements only in one portion of the underground mine, consisting of the ramp that connects the ore zones to the surface, i.e. a tunnel narrow enough to allow the presence of only one truck and featuring occasional passing bays where trucks can pull into in order to allow equipment or trucks travelling in the opposite direction to pass. The rest of the mine, consisting of the various underground levels, the haulage networks in each level, the loading points in each level etc., is neglected, and approximated in the model with a single loading point, capable of serving one truck at a time. Another limiting drawback of the model is that it considers all vehicles to be located on the surface at the beginning of the shift, contrary to what happens in most underground mines. The tool, furthermore, makes plans for the vehicle movement for an entire shift. By launching the optimization model at the beginning of the shift, the schedule of each vehicle for the entire working shift is found. In articles [6,14,15] we find justified criticism of the use of such long-term planning techniques in the underground mining context. Another shortcoming is that the number of trips performed in a shift is a given input of the model. The authors, however, presented a way to circumvent this limit. Finally, [24] considers deterministic loading, unloading and transportation times. To conclude, all the disadvantages already discussed for the other articles apply.

Still in the field of fleet management in underground mines, [18] is one of the few works that attempt to take into account the stochasticity of the underground world.

The presented approach, however, has the drawback of not considering the entire underground haulage network, but only a ramp in which traffic lights manage access to the one-lane segments. The reader should note that, when considering only a ramp controlled by traffic lights, the fleet management problem requires fewer decisions to be taken. The routing part of the problem disappears, as for each pair of departure-destination points the road is unique. Traffic conflicts don't have to be considered implicitly in the schedules to be created for vehicles anymore, and the scheduling decision is reduced to the control of traffic lights. The only problem that remains the same, therefore, is the dispatching.

In this type of context, [18] proposes a cooperative coevolutionary algorithm able to simultaneously generate the traffic light schedule and the truck dispatching plan for a fixed future time window. The truck dispatching plan consists of an ordered list of shovels (representing loading points and, therefore, destinations); each vehicle that will arrive on the surface in the time horizon considered will be dispatched according to this list. It should be noted that in [18]'s model there is only one dumping point at the surface, and, therefore, there is no dispatching decision for vehicles departing from loading points. The traffic light schedule, instead, consists, for each traffic light in the ramp, of a list containing time intervals in which the lights will be green and red respectively. One of the following KPIs is used during the algorithm run: total trucks waiting time; average trucks cycle time. These short-term KPIs are chosen since considered proxy metrics of the overall long-term objective, embodied by the throughput of the transportation system (measured in total truckloads unloaded per shift). The pair of schedules (one for the dispatching, and one for traffic lights) that minimizes the selected KPI, compared to all the pairs generated by the algorithm, will be the output of the procedure. Simulation is used to evaluate the performance of the schedules.

The used simulator is based on a network of timed automata. For more information on the type of simulation employed and the concept of timed automata, the reader can refer to [15]. The important aspect to remark, however, is that the fill, empty and travel times considered are not deterministic values, but randomly generated from a uniform distribution. In this way the proposed tool models the variations that occur in the real world mine and takes into account the stochastic aspect of the problem. The algorithm, to summarize, is launched at a given time to produce a temporary, short-term schedule for a fixed future time window. "Periodically, this schedule is discarded and the Evolutionary Algorithm is rerun using updated information to ensure that the schedule remains useful over time under variable real-time conditions" [18]. The statement just quoted embodies the strength of [18]'s approach with respect to the articles considered earlier in this chapter. The author, in fact, is aware that the plan for the future made in a certain moment of time will become obsolete in the future itself, due to the variable and unpredictable conditions that constitute the nature of an underground mine.

The fleet management technique presented by [18], however, would be difficult to implement for the entire mine, since it considers only a ramp managed with traffic lights. Among the various reasons, there are: the incapacity to manage sections of the transportation network having a different nature from that of the ramp managed by traffic lights; the failure to take into account the routing aspect of the problem; the excessive computational effort of the algorithm, which conflicts with the need to take real-time decisions with a few seconds to spare (according to [25], in fact, in most real-life applications of evolutionary algorithms the computational aspect is prohibitive).

[19] proposed a decision support instrument for optimizing the schedule of the entire mining process, applicable in the case of underground hard rock mining methods.

The tool, taking as input the short-term plan and the real-time conditions of the mine (considering both underground equipment and working faces), is able to create a schedule for the entire shift, offered to the decision maker in the form of a Gantt chart showing what time which specific machine (e.g. driller, explosive charger, bolter) should operate on which specific working face. This timetable is obtained by minimizing the timespan to perform the specific workload foreseen in the short-term plan.

A mine may use the tool, as claimed by the authors in their study, to create the schedule for resources and/or to quickly reschedule vehicles when deviations from the plan programmed by the instrument occur, for example due to unexpected events.

[19] presented their study as a means to fill a gap existing until then in the underground mining literature, namely the lack of consideration of the overall underground mining process. The two main mining processes carried out in parallel in an underground mine, more specifically, are the development and the production. Each of the two processes requires that the working face in question is visited sequentially by a series of mining vehicles, each aiming to carry out a particular operation (e.g. drilling, explosive charging, bolting). The plus point of [19]'s approach is precisely that of not focusing solely and exclusively on the loader-truck cycle as done by all the other studies in the literature, but of having a more global perspective that considers the entire mining processes, including all the individual activities embedded in them.

The proposed solution, however, has the crucial drawbacks of creating schedules on a shift-to-shift basis (philosophy negatively criticized in [6, 14, 15]) and of addressing only the dispatching aspect of the problem, neglecting the routing and the scheduling.

What is more, [19], like the studies already reviewed, has the disadvantage of considering deterministic input data. The negative consequences of doing so are further amplified in [19], since plans are created for entire work shifts. This drawback, however, is to be counterbalanced with the capability of the tool to provide a solution in a short time, as claimed by the authors in their article: whenever actual operations deviate from the timetable foreseen by the tool, it is possible to rerun the procedure and have a new schedule in a short time. The schedules and re-schedules obtained in this way, however, although generated with a high reactivity, give rise to a timetable far from optimal. In a nutshell, the inventors of the instrument are aware of the need to modify the proposed solution in order to better consider the high uncertainty in underground mining. Also [9] negatively criticized the instrument, remarking that it does not include unforeseen events and uncertainty in execution.

The sole approach in the literature, to the best of the author's knowledge, which is neither myopic, nor deterministic, nor limited to the transportation fleet, is [26].

[26]'s intuition consists in using dynamic scheduling and discrete event control to provide planning, dynamic control and coordination of automated vehicles in an underground mine.

As far as the dispatching aspect of the problem is concerned, the task assignment for a vehicle that has completed the current operation and asks for a new destination is performed by a dynamic scheduler, which considers all the other vehicles in the network, their positions, their goal destinations, and their planned routes.

With regard to routing, given the current position and a goal position, simple path-finding routines (such as a depth first or breadth first search) used in combination with discrete event control are able to find a path for the vehicle in question. Roads already planned for other vehicles can, if necessary, be questioned and modified.

The discrete event control, its related control structure, and its related controller, finally take care of the scheduling part of the problem. The destination and the road assigned to a vehicle in the previous steps are translated into sequences of elementary micro-activities to be performed (e.g. movement from the current position to the end of the road segment), modelled using Petri net formalism. The discrete event controller regulates the traffic in the mine by controlling the firing of transitions in the Petri net structures of the different vehicles, so as to have coordination (e.g. regulation of access to a shared resource, such as an intersection). A vehicle, therefore, will reach its final destination by performing unitary movements and coordinating with the others thanks to a continuous exchange of messages with the controller.

The result of integrating dynamic scheduling and discrete event control is a system with a high degree of reactivity. The vehicle is only informed about the current movement, while the future path and destination can be modified at any time. Whenever a task has to be assigned to a requesting vehicle, there is a traffic conflict, or there is an unforeseen event, the road and/or destination assigned to each vehicle can be changed. The whole system, therefore, could be re-planned at any time, adapting to the real-time conditions of the duo system-environment.

The FMS presented in [26], giving life to a system with a high degree of reactivity, embraces the vision and philosophy of this research. This is why the methodology of this thesis and the one used in [26] have some common points (such as the employment of dynamic scheduling).

[26], however, places its approach on a principally theoretical level, and there are clear shortfalls in the methodology, especially with regard to the dispatching and the routing. The main

purpose of the paper, in fact, is to present to the reader a general theoretic framework and a control architecture that integrate discrete event control and dynamic scheduling to solve the fleet management problem.

Another journal article introducing dispatching, routing and scheduling strategies is [20]. The object of the study, however, although still an underground mine, is radically different from the one considered in this research, because the underground mining method implemented is the block cave mining technique. Although the authors of the article claim that their strategies are not limited to this specific mining technique, the layout of the mine and the logic of movement of the various fleet vehicles are different from those dealt with throughout this research. The three fleet management problems, i.e. dispatching, routing, and scheduling, are, therefore, significantly different. The management strategies presented by [20], as consequence, will not be further investigated.

2.3 Pragmatic analysis of academic contributions

Table 2.2 presents the reader with an overview of the weaknesses of the FMSs proposed in the literature and discussed in Section 2.2.

Table 2.2 Shortcomings of the FMSs proposed in the literature

	Deterministic approach	Myopic approach	Approach limited to the transportation fleet	Approach failing in the management of unexpected events
[6]	X	X	X	X
[16]	X	X	X	X
[17]	X	X	X	X
[24]	X	X	X	X
[18]		NA[1]	X	
[19]	X	NA		
[26]				

LEGEND

[1] *NA stands for non-applicable*

The performed analysis shows that flaws very frequently encountered in the literature solutions are: the lack of consideration of the stochastic nature of the system variables (the approach, in this case, is called deterministic); the lack of consideration, when making decisions for a vehicle, of the vehicles that will have to be managed in the future (the approach in this case is myopic and egoistic); the lack of consideration of the entire development cycle (focus only on transportation and transport vehicles); the lack of a method for managing unexpected events. The first and last shortcomings can be combined into one, i.e. the failure to take operational uncertainty into account.

In this section, always referring to the literature, the criticality of the drawbacks just listed will be highlighted. The latter leads to the need for FMSs with certain characteristics (summarised in Table 2.3). Understanding what characteristics a FMS must have in order not to pose the problems outlined in the table is of crucial importance, especially in the design phase of a FMS. Since the conception of a FMS is one of the objectives of this research, this paragraph is of paramount significance to understand the choices that will be made in the rest of this thesis. In a nutshell, the strategy is to study the weaknesses of the FMSs presented in the literature in such a way as to conceive a FMS able, as far as possible, to avoid them.

Table 2.3 Translation of flaws in literature FMSs into requisites for an ideal FMS

FMS shortcoming	Consequences	Indicators positively correlated to the gravity of the consequences	Features of an ideal FMS
Non consideration of operational uncertainty (stochasticity, unexpected events)	Deviations; utopian plans; traffic conflicts; sub-optimal solutions; resumption of decisions already taken in the past; etc.	Randomness of system variables; frequency and magnitude of unexpected events	Consideration of operational uncertainty in the planning phase
Myopicity	Worsening of the mining system KPIs with respect to a collaborative approach		Non myopicity, collaboration
Focus only on the transportation part of the problem, and not on the whole development cycle	Solution related to a simplified problem, difficult to implement in reality		Focus on the whole development cycle, and not only on trasportation

As explained also in [27], creating plans for the future without considering operational uncertainty results in deviations between the plan created and its execution. This can mean trying to implement in reality plans whose objectives are difficult or even impossible to achieve. More specific details on the negative consequences of the use of deterministic approaches have already been presented during the analysis of [6]; for the record, among them are: traffic conflicts, sub-optimal plans, need to revisit decisions already taken in the past. All these problems resulting from the lack of consideration of operational uncertainty are all the more acute the greater the aleatory nature of the system, and the greater the frequency and magnitude of unexpected events. In the 85% of the articles reviewed throughout the entire course of this literature review (twins, thus, of this thesis in terms of topics covered) there is at least one sentence in which reference is made to the high aleatoriety, operational uncertainty, and/or unexpected events (or, in any case, to concepts related to the three listed above). This statistic makes the reader understand that the high level of operational uncertainty is an intrinsic characteristic of the nature of any underground mine, and highlights the magnitude of the problems arising from its lack of consideration during the planning phase. In the design phase of a FMS, therefore, the inventor must ensure that the developed tool considers operational uncertainty.

The importance for a FMS to consider not only the vehicle that requires a new destination, but also those that will do so in the future, is explained in [6, 14, 26]. The KPIs related to the performance of the mining system would undoubtedly be superior if a FMS with this characteristic were adopted.

Last but not least, the last trait that a good FMS should have is the consideration of the entire development cycle and the entire vehicle fleet, and not just the transport activity and the transport vehicles (this requirement was pointed out also in [19], just to mention one). Otherwise, as happens in most of the revised FMSs, the obtained solution would relate to a simplified problem, far removed from the real problem faced by underground mines, and it would therefore be difficult to implement in reality.

In order to conceive a FMS that does not suffer from the typical shortcomings of existing FMSs, and therefore able to dodge the related problems, one of the possible ways is to consider the ideal features shown in Table 2.3 during the design phase of dispatching, routing and scheduling strategies. This direction has been taken in the past only by [26], already examined in Section 2.2, and will be the one followed in the course of this thesis.

As concerns the consideration of the operational uncertainty, however, it should be remarked that solutions other than the tailored design of dispatching, routing, and scheduling strategies have been proposed in the underground mining literature. An example is [27], which manages to consider operational uncertainty by introducing a simulation phase into the planning process. The proposed simulation-optimization framework, in particular, is able to generate short-term plans (indicating when and where to carry out which activities) characterized by a high probability of being able to be effectively executed in reality, thanks to the consideration of operational uncertainty. The optimization, able to generate short-term plans considering deterministic data, is integrated with the simulation, used to take into account the operational uncertainties of the mine. Uncertainty, in fact, would be difficult and cumbersome to incorporate into a mathematical optimization model. The proposed solution, more specifically, consists of an iterative calculation, in which in each iteration first a short-term schedule is generated using a mixed-integer linear programming model, and then the latter is simulated using a discrete-event simulation model. The output of the simulation is then used as input for the optimization, starting the next iteration. After each iteration there is an increase in the probability that the obtained plan is actually implementable in reality (defined as adherence). When a certain adherence threshold or time limit is exceeded, the calculation stops and the last plan created is considered. As uncertainty has been considered during the planning phase, with this approach the negative consequences in Table 2.3 are avoided.

Revising [27] was useful, moreover, to exclude stochastic programming from the portfolio of techniques that the author could have exploited to shape his own FMS. Stochastic programming, in fact, being capable of modeling optimization problems that involve uncertainty, would have been a suitable approach to consider uncertainty in the planning phase (feature of an ideal FMS, as explained in Table 2.3). In the planning context of an underground mine, however, it would be very difficult, and most likely impractical, to implement such a technique [27]. The possibility of applying this approach was therefore discarded.

2.4 State of the art - Commercial software

The focus of this research project is the fleet management problem in underground mines. In order to expose the full spectrum of solutions currently existing for this problem, in parallel to academic models, algorithms or, more generally, approaches, commercial FMSs available on the market should be considered. The dispatching, routing and scheduling problems discussed by this research, in fact, are also addressed by some commercially available mining solutions offered by companies such as Siemens AG, ABB, Hexagon Minning, Cisco Systems Inc., AT&T Inc., IBM, Verizon Connect, Geotab, Masternaut and Trimble. Unfortunately, the underlying logic and algorithms implemented in these industrial software packages are proprietary information not disclosed into the public domain. Due to their proprietary nature, these industrial FMSs cannot therefore be further explored, and their level of analysis unfortunately stops here. Considering both underground and open-pit mines, the only disclosed commercial FMS is DISPATCH®. In the 1980s and early 1990s, in fact, [28] and [29] revealed in the public domain some models and algorithms for fleet management, which were used some years later by Modular Mining System to develop the DISPATCH® FMS. DISPATCH®, by the way, is a solution suitable only for open-pit mines. Given the non-transferability of fleet management solutions for open-pit mines in underground mines [14–17], this commercial software will not be further investigated throughout this section.

CHAPTER 3 RESEARCH BACKGROUND

3.1 Introduction

This chapter will first highlight the scientific contributions, then the research objectives, and finally the theoretical concepts underlying this work.

3.2 Research contribution

The objective of this section is to introduce to the reader the approach used in this thesis to solve the fleet management problem in an underground mine, to demonstrate its theoretical validity, and to locate it with respect to the literature. First of all, the main concepts behind the proposed fleet management solution will be presented. Once this is done, literature will be employed to validate this solution from a theoretical point of view. Subsequently, it will be demonstrated that the proposed approach is innovative, due to its capability to fill a gap in the literature of underground mines, and at the same time promising. The proposed FMS will be compared, finally, to those presented in the literature: it will be pointed out that it is free from the typical shortcomings of the latter, and, consequently, able not to bear the related problems.

The solution for the fleet management problem in underground mines suggested by this thesis consists of a distributed approach based on *agent-based systems* and *dynamic multi-agent scheduling*. The idea is to implement a distributed mining system made up of software agents, in which each vehicle is autonomous and able to take dispatching, routing and scheduling decisions on its own, communicating and coordinating with the other members of the fleet. The management and control system thereby created is no longer centralised, as happens in the FMSs of the literature, but distributed and decentralised.

In the light of the widespread use of agent-based technology in traffic and transportation systems [30], a FMS founded on agent-based approaches seems to be suitable for the fleet management in an underground mine. However, due to the fact that such a solution has never been proposed until now, we looked for clues in the literature in order to understand whether the proposed solution could actually be implemented (at least from a theoretical point of view) to efficiently and effectively manage the fleet of an underground mine.
[14], [6], [16], [17], and [20] highlight the similarity of fleet management problems in underground mines and in the field of AGVs. Considering this, at least for the moment, as a hypothesis, we can introduce the following thesis: the techniques used in AGVs FMSs can

be used, at least from a theoretical point of view, in underground mining FMSs. In support of this thesis there is [6], which, endorsing the above hypothesis, used a solution adopted in the field of AGVs to propose its own FMS for underground mines (obviously adapting the solution to the different context, in order to make it valid, more than just theoretically, also practically). Second evidence underpinning the thesis is the fact that the literature reviews of [6,14,16,17,20], which are all articles proposing FMSs for underground mines and supporting the above hypothesis, are full of references to works on AGVs: this means that also for all these authors the solutions adopted for AGVs are transferable to the mining world. These two proofs are considered sufficient to prove the validity of the previously stated thesis on the transferability of solutions from the world of AGVs to that of underground mines, given the hypothesis of equality of the two fleet management problems. There remains nothing else to do, therefore, but to demonstrate the truthfulness of this hypothesis. We can claim that for AGVs and underground mines the same fleet management problem occurs, due to the equality between: (1) the systems to be managed; (2) the transportation networks in which these systems operate; (3) the decisions to be taken to manage these systems, i.e. the objectives of the fleet management systems.

(1) The system to be managed in both contexts is represented by a fleet of vehicles. The goal of an AGV is to load, transport and unload materials from one area to another of an industrial building. Load-haul-dump vehicles and trucks perform exactly the same task inside a mine. Vehicles such as drillers, chargers, bolters and cleaners aim to perform certain activities at specific points in the underground mine; even if we are talking about activities other than loading/unloading materials, the movement of the vehicle from one point to another of the mine remains the main operation to be managed. What AGVs and all mining vehicles have in common, therefore, is the definition of the vehicle operating cycle: given the current position, corresponding to the destination assigned in the previous operating cycle, the vehicle will travel up to a new destination, where it will perform a certain activity concluding the operating cycle in question. What changes according to the nature of the vehicle is the meaning of destination point and the activity to be carried out in it. But what the management system must be able to handle, taking dispatching, routing, and scheduling decisions, is the operating cycle, defined in the same way in the contexts of AGVs and mines.

(2) Underground haulage networks are composed of one-lane bi-directional road segments (i.e. they can contain a maximum of one vehicle in width, and the traffic flow is allowed in both directions), and intersections are not large enough to allow the presence of more than one vehicle at a time. The same traits can be found in many of the networks used for AGVs navigation. Evidence of the equality of the transportation networks in the two contexts can also be found in [6,14,16,17,20].

(3) As already mentioned, the fleet management system aims to govern the vehicle operating cycles, regardless of whether we are talking about AGVs or mining equipment. The decisions to be taken, more specifically, are the same in both contexts: dispatching (assigning a destination to the vehicle); routing (determining which path to follow to reach the destination); scheduling (controlling traffic in the network to avoid conflicts).

In the light of the above, we can assume to be true both the hypothesis of equality of the two fleet management problems, and the thesis of transferability of the techniques used to manage AGVs in the underground mining world.

In more pragmatic terms, we will now make use of these findings to demonstrate the theoretical validity of the solution proposed by this thesis. We can start by observing that this research presents a solution based on the agent-based paradigm, and that in [31–37] such agent-based approaches are deployed to manage a fleet of AGVs. But then, since the solutions used for AGVs fleet management are transferable to the underground mining world, we can say, even before the actual FMS is designed, that it, at least from a literature-based theoretical point of view, can actually be implemented in underground mines to efficiently and effectively solve the fleet management problem.

During an in-depth analysis of the literature, no work in which agent-based systems are used for the fleet management problem in underground mines was found. Taking a broader perspective and considering the whole mining sector, [38] was the first to use an agent-based solution for the fleet management problem, leveraging on this paradigm to create a FMS valid for open-pit mines. In [38], in fact, the authors declared that “to the best of their knowledge, there was no evidence that before them a multi-agent system had ever been applied to truck dispatching in open-pit mines”. Due to the lower complexity of the problem, we know that many more solution approaches have been proposed for open-pit mining than underground mining; evidence of this gap between the two contexts can be found in [6, 14–17, 19, 20]. We can deduce, therefore, that there is a high probability that the agent-based paradigm, used in open-pit mining for the first time in 2020, has never been proposed for the fleet management problem in an underground mine. [38]’s assertion, therefore, supports the author’s search results. In a nutshell, the conclusion is that, to the best of the author’s knowledge, nobody in the past has ever used agent-based systems to propose solutions to the fleet management problem in underground mining. This thesis, therefore, is intended as a way to fill this gap in the literature, contributing to development and innovation in the field of operations management in underground mines.

Besides proposing something never presented in the past, this thesis seems to be able to obtain positive results. A first clue is given by [30], which, conducting a literature survey on agent-based technology applied in traffic and transportation systems, demonstrated the high potential of this tool to improve the main performance of these systems. Secondly, in the light of the demonstrated equality of fleet management problems in underground mines and AGVs, and knowing that agent-based systems are widely and profitably employed in the field of AGVs (renowned for their ability to lead to cost effective movement of materials), we expect this thesis to propose a FMS capable of efficiently and effectively manage the fleet of mining vehicles. Likewise, also the results obtained by [38], which proposed and evaluated through agent-based simulation an agent-based system for open-pit mines capable of achieving an overall improvement of the KPIs of the mining system, highlight the potential of this research project.

During the design phase of the FMS presented in this thesis, then, the shortcomings of existing FMSs were taken into account, with the aim of developing a solution capable, as far as possible, of avoiding them and not posing the related problems (Table 2.3). The agent-based approach has proved to be very useful to achieve this objective: leveraging the advantages offered by this paradigm, such as the capability to deal with uncertainty and dynamism, and the capacity to efficiently and effectively synchronize activities, the proposed FMS is free from the weaknesses and problems of the solutions presented in the literature, as one can realize by reading the rest of this thesis.

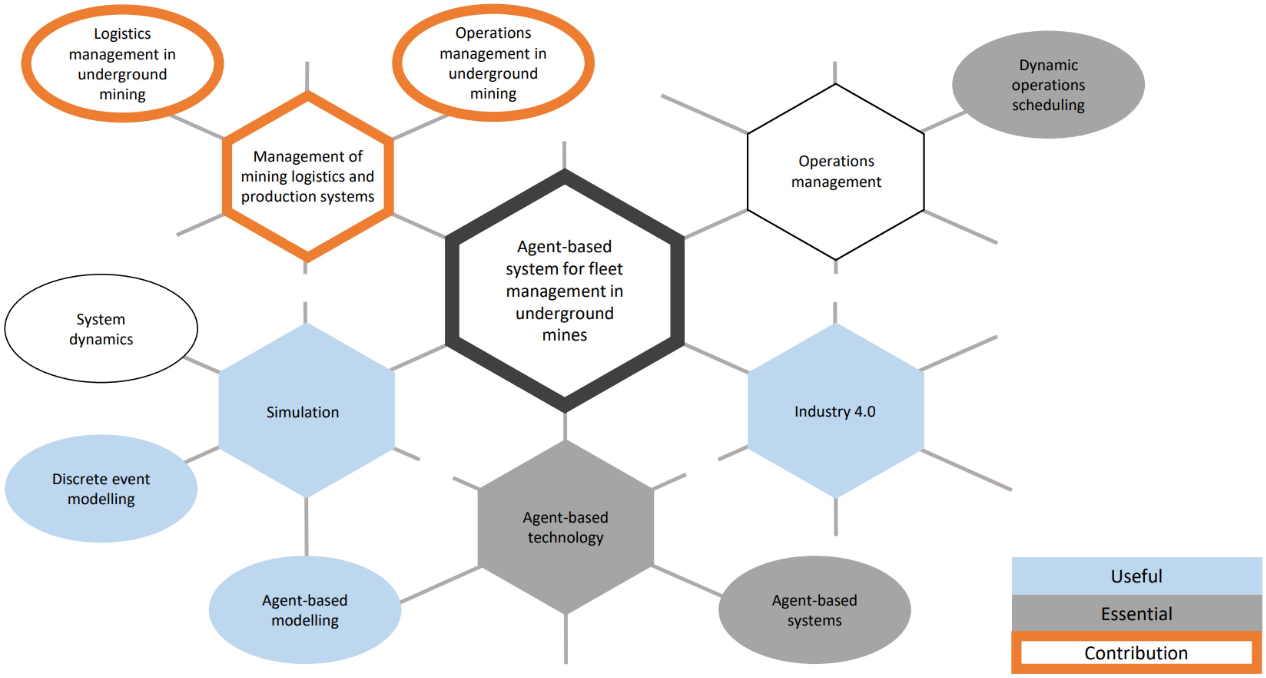

Figure 3.1 Areas of Relevance and Contribution Diagram [1]

Figure 3.1, in conclusion, portrays the Areas of Relevance and Contribution Diagram [1] of this research project, highlighting the fields which are essential and useful for the project, and those in which the research develops new knowledge.

3.3 Research objectives

Chapter 1 has contextualised and clarified the meaning of fleet management in an underground mine. Chapter 2, on the other hand, has highlighted a knowledge gap in the literature, which justifies the research objectives that are about to be presented in the following section.

With a thorough picture of the problem to be addressed in mind, and being aware of how well agent-based systems[1] are able to solve the exact same problem in the industrial context of automated guided vehicles [31–37], the author wonders whether such a solution is transferable to the world of underground mining. The research question connected to this idea is:

RQ: Can agent-based systems and dynamic multi-agent scheduling solve the fleet management problem in an underground mine, leading to an efficient and effective logistics and production system?

The whole thesis is structured around this central question, which places agent-based technology as the cornerstone of the proposed solution to the fleet management problem. Once established the “why” of the research, the author investigates the “how”, and comes up with the following research objective:

RO: This research aims to design and assess the performance of an agent-based system and a dynamic multi-agent scheduling procedure capable of managing the fleet in an underground mine.

The research objective is founded on the merger of two key aspects, which support one another to form an integrated whole: the design and the evaluation of a solution to the fleet management problem. The single aspects underpinning the research objective are then disentangled from it, giving rise to the following research sub-objectives:

RSO1: Design an agent-based system and a dynamic multi-agent scheduling procedure for the management of the fleet involved in the extraction and transportation processes of an underground mine.

RSO2: Build a digital replica of a real underground mine (industrial partner), and use simulation to reproduce the behaviour of the designed agent-based system and dynamic multi-agent

---

[1] *Agent-based systems* will also be referred to as *multi-agent systems* in the course of this thesis: one term is preferred over the other depending on whether one is emphasizing the agent-based paradigm or the community of interacting agents.

scheduling procedure, in order to quantitatively assess their performance.

RSO3: Conduct a quantitative benchmarking between the proposed solution to the fleet management problem and the one currently adopted by the mine under review.

RSO4: Conduct a qualitative benchmarking between the proposed solution to the fleet management problem, the one currently adopted by the mine under review, and the principal ones presented in the underground mining literature.

The first sub-objective (RSO1) is the centrepiece of the research: an innovative solution to the fleet management problem in underground mines will be devised. The remaining three sub-objectives play the role of evaluation. Simulation modelling, in particular, will be employed as a decision support tool to answer the research question, i.e. to understand whether the proposed solution is actually able to solve the fleet management problem in an efficient and effective manner. By implementing the agent-based system and the dynamic multi-agent scheduling approach in the digital replica of the industrial partner, it will in fact be possible to evaluate their performance, assessing the impact on the key performance indicators of the mine (RSO2). Last but not least, the conceived fleet management system will be evaluated through comparison with the one adopted by the industrial partner (again thanks to the created simulation model) and the major ones proposed in the literature (RSO3 and RSO4).

With the purpose and the mission of the research project in mind, the Initial Reference Model [1] in Figure 3.2 has been designed. The latter provides a conceptual framework, which defines the variables relevant to the study and maps out how they might relate to each other. In a nutshell, it reflects what the author expects to find through this research.

All terms found in the Initial Reference Model have already been introduced in the present and previous chapters. It is worth noting that items “myopia of the fleet management system”, “fleet management system limited to transportation” and “neglect of uncertainty” are weaknesses in the solutions proposed for the fleet management problem in the underground mining literature, and have been cited by the authors themselves. As this research aims to devise a fleet management system, during the design phase the author considered the shortcomings of the systems already present in the literature, in order to propose, if possible, a solution exempt from them.

After this introductory digression, the meaning of the Initial Reference Model can now be explained. It consists, in particular, of two parts. The former is the exogenous part of the model and concerns the constraints of the fleet management problem, already unveiled in Chapter 1. Regardless of the quality of the management system (endogenous part of the problem), the presence of the four constraints shown in Figure 3.2 (in the bottom left-hand

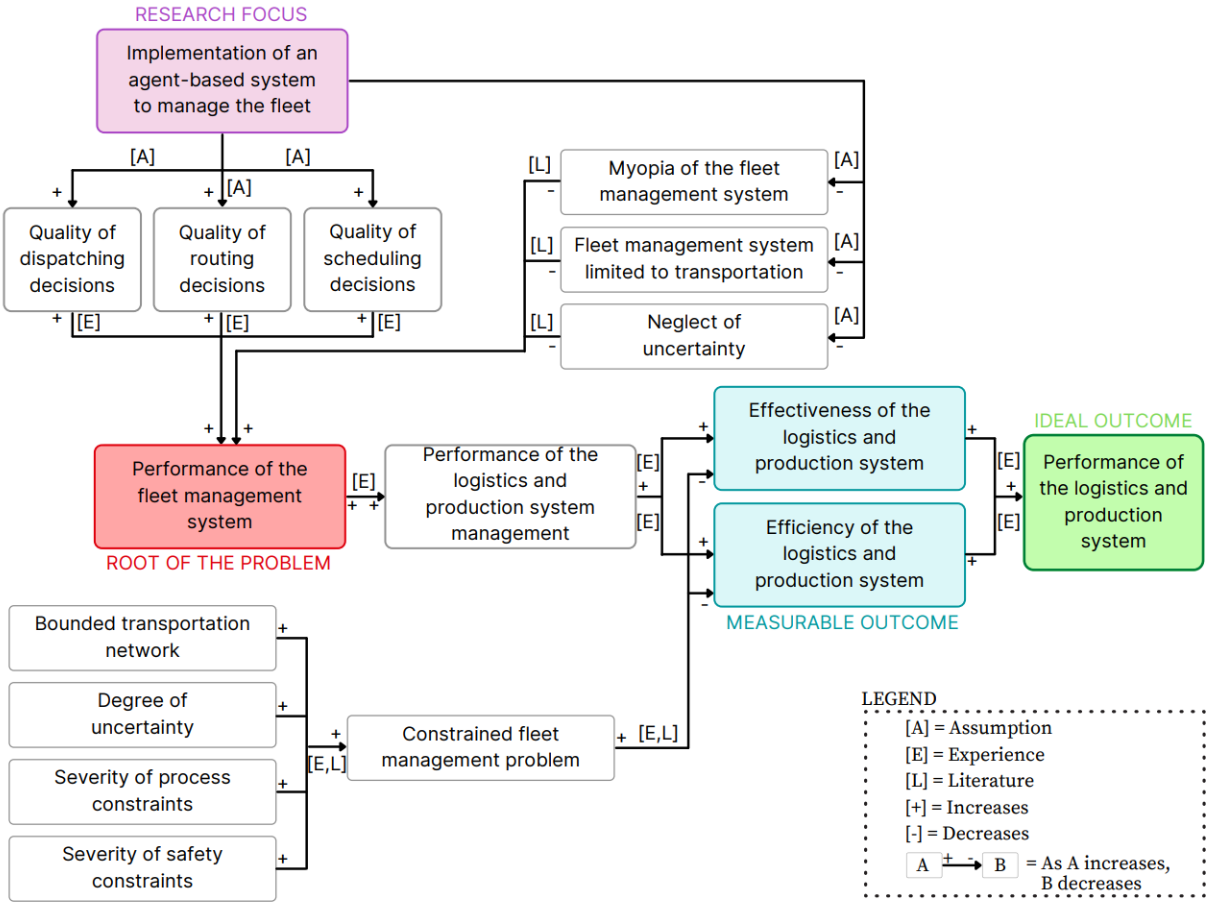

Figure 3.2 Initial Reference Model

corner) hinders the mining system in achieving its objectives of efficiency and effectiveness. The latter is the endogenous part of the model, and echoes the research question presented above. The assumption is that the employment of a tailor-made agent-based system leads to: better management decisions (dispatching, routing and scheduling decisions pinpointed in Chapter 1); consideration of uncertainty; a management system that is neither myopic nor limited to transportation. All these elements contribute to the quality of the management of the fleet. But fleet management is only a slice of the overall management of the logistics and production system. Improving fleet management, therefore, means improving the comprehensive management of the mining system. But better management translates into increased efficiency and effectiveness of the managed system. Hence, to close the circle, if the opening assumption were true, the initial agent-based system would lead to a more efficient and effective logistics and production system. This is exactly what was inquired into in the research question. The aim of the following chapters, to conclude, will be precisely to test the veracity of the assumption, and to find an answer to the research question.

3.4 Theoretical framework

Chapter 1 outlined the profile of the problem addressed by this research, namely the fleet management problem in underground mines. The research mission and purpose presented in Section 3.3, on the other hand, hinted at the solution that will be proposed in the rest of the thesis, i.e. an agent-based system, and the method that will be used to evaluate it, i.e. the simulation. Given the central role of these two aspects in the research project, the concepts and theories underlying them will be further explored in this section. When diving into the agent-based world and talking about agents and agent-based models, there is no standard definition that is universally agreed upon. This chapter therefore also aims to align the reader with the author's viewpoint, by pinpointing the definitions in the literature taken as reference.

3.4.1 Simulation modelling

One of the main objectives of this research project is to create a simulation model of a real underground mine (industrial partner). Hence the following section will elaborate on the meaning of simulation modelling.

The term "simulation" refers to a set of principles and techniques allowing to reproduce the functioning of a real-world process/system. A simulation study is generally coupled with a decision problem (e.g. improvement or design of a system), and this makes simulation a decision-support tool for evaluating the performance of a solution to a problem. It is usually performed by means of computers, with the aim of creating a digital representation of the studied system and conducting experiments in a risk-free world. The term "simulation modelling", in fact, indicates that what is simulated is a model of the real system, obtained by mapping the real world to the world of models by choosing the desired level of abstraction and the proper modelling language, and conducting all the other practices that constitute the art of modelling. The modeller, therefore, seeks the solution to the problem by immersing himself/herself in a virtual, risk-free world in which it is permissible to make mistakes, learn from them, go back and start again [2].

In order to map a real-world system to its model, modern simulation modelling uses three methods, each serving a specific range of abstraction levels: discrete event modelling, agent based modelling, and system dynamics [2]. Figure 3.3, in this respect, has been adapted from [2], and shows the level of abstraction accessible with each modelling method.

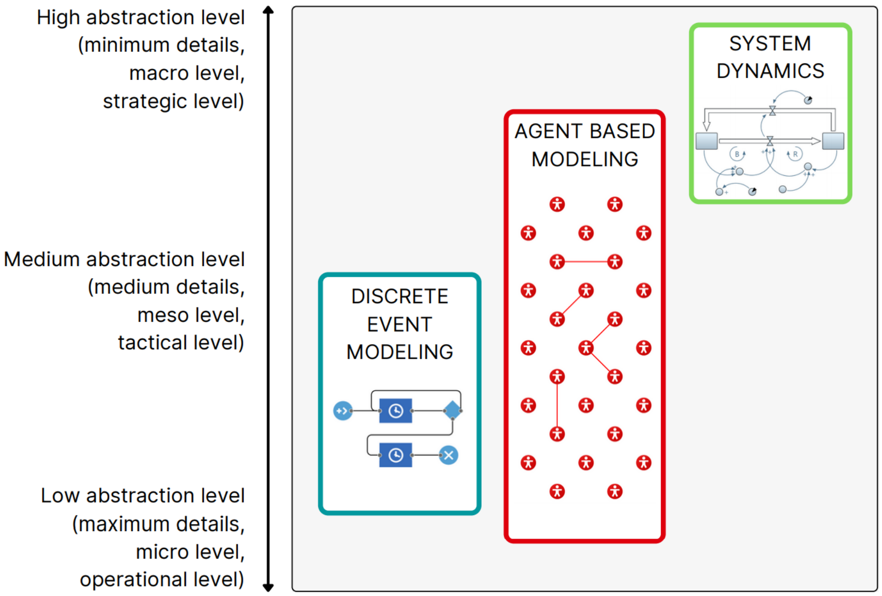

Figure 3.3 The three methods in simulation modelling, by [2]

The choice of the right method is an art rather than a science and depends on the system to be studied and, above all, on the objectives of the study. Figure 3.4, for example, shows how the same real-world system (an underground mine) can be modelled using different methods depending on the goals of the study. Let us suppose, for example, that the objective is to study the management of operations in the above-ground plant, to which the extracted ore is brought by the logistics system. In this case, the modeller could employ discrete event modelling, creating process flowcharts in which units of raw material (extracted ore) would become the entities, and manufacturing machines and employees the resources. If the internal logistics was being investigated, instead, agent-based modelling could be used, turning vehicles into agents that exert mutual influence while travelling in the bounded transportation network, and that are influenced by the tasks assigned by the mine engineers. Or, finally, if the study in question concerned maintenance strategies, and the company was interested in studying the relationship between preventive maintenance frequency, mean time between failures, maintenance costs and fleet availability, systems dynamics would be the best modelling tool.

This thesis will make use of multi-method simulation (combining agent-based and discrete event simulation), with agent-based simulation being predominantly present. The reasons behind this choice and evidence of how the adopted approach is consistent with both the studied system and the research objectives will be better illustrated in the methodology section (Chapter 4).

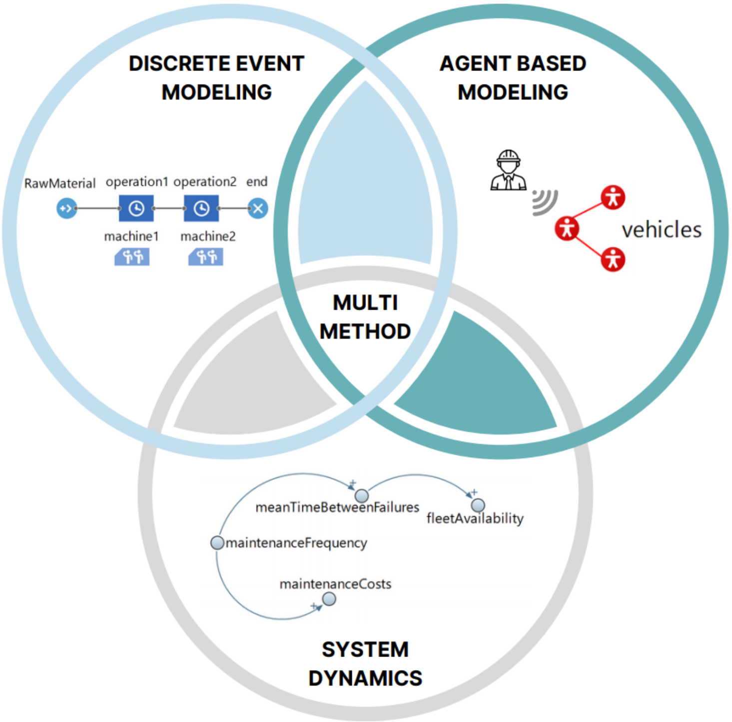

Figure 3.4 Choice of the modelling method best suited to the objective of the study

To conclude the digression on simulation modelling, it should be remarked that a simulation study offers several advantages. Among them are: the possibility of analysing a system without perturbing it; the possibility of studying the link between the internal functioning and external behaviour of a real system, with a view to improving it; the possibility of anticipating the performance of a system without having to build it physically in the real world. This research project, in this regard, will exploit all the above-mentioned simulation benefits. The logistics and production system of an existing gold mine will be analysed (without any disturbance) by investigating the relationship between the behaviour of individual resources and the overall behaviour of the mining system. The performance of a multi-agent system capable of managing the fleet will then be evaluated through simulation, without the need to implement such a solution in the real-world mine.

3.4.2 The agent-based world

This section will introduce *agent-based models*, *agent-based simulation*, and *agent-based systems*, as they are the pillars of this thesis. First, however, the concepts of *distributed system*, *agent* and *environment* will be reviewed, as they are propaedeutic notions.

Distributed systems

Mining systems, like almost all logistics and production systems, are distributed. In a distributed system, there are: several actors in an environment over which they have a varying extent of control; decisions of different nature taken in parallel, which lead to parallel actions; interdependencies and interactions between the actors, within and through the environment. Bottom line, the behaviour of such a system is not centrally controlled, but it emerges from the actions and interactions between these unitary entities. It is easy to realise that most of the systems around us are distributed: supply chains, traffic networks, hospitals, factory warehouses, etc., all feature the above-mentioned characteristics.

In a mining logistics and production system, vehicles and their drivers are the main actors living in a shared environment, taking decisions and acting in parallel, and exerting a mutual influence on each other. The behaviours and interactions of these players bring out the behaviour of the overall system, which cannot therefore be defined *apriori* and in a central manner. The nature of the mining logistics and production system is therefore distributed by definition.

Being aware of the nature of the studied system is of paramount importance in order to identify the right tools for analysing and re-designing it. In the case of this research, the distributed nature of the mining system will justify the use of agent-based modelling and simulation on the one hand, and the employment of agent-based systems on the other, as tools for analysis and design respectively. Agents in agent-based models/systems, in fact, are the perfect twins of actors in distributed systems, as will be explained later.

Notion of agent

The concept of agent is the cornerstone underpinning this research project. It takes on different nuances depending on the context: when conceiving a model of a real system or simulating it, an agent can be seen as a modelling unit; when, instead, we consider agent-based systems (real physical systems made up of interacting elementary software systems, i.e. the software agents), with the term "agent" we refer to a computer programme built from a detailed model of the agent, which executes the model of the agent. Regardless of the reference context, however, an agent is characterised by the fact of being situated in an environment, from which it perceives information, and on which it acts. [39] in fact defined the agent as an entity that: *"operates in an environment from which it is clearly separated; makes observations about its environment; [...] initiates and executes actions to change the environment"*. Moreover, there is no universal agreement in the literature on the precise definition of agent beyond the essential property of autonomy. The definition adopted by this thesis is based on the one proposed by [40]. According to the author three main characteristics are necessary in order to define an agent as such: it must possess a certain degree of autonomy; it must possess a behaviour; its state (set of variables characterising it) must be able to change dynamically either as a result of its own actions or as a result of the interaction with its environment. Additional properties are interaction with other agents and action planning. The behaviour and actions of an agent may be guided by simple rules of behaviour (knowledge), by objectives, and by learning from its past behaviour and actions (intelligence) [41]. An agent, to conclude, can be reactive or autonomous, solitary or social, selfish or collaborative, rigid or agile, non-adaptive or intelligent. The meaning of these attributes and the concept of agent will be explored in more depth throughout this thesis.

Notion of environment

The notion of environment can be studied from a dual perspective. If we focus on the system-environment dichotomy, the environment is nothing more than the dynamic context in which the system evolves. If, on the other hand, we focus on the agent-environment dichotomy, the environment represents the set of elements that the agent perceives outside itself. In the latter case, the concept of environment is tied to that of information, which is why it is often referred to as “informational environment” of an agent. The concept is also relative: each agent has its own environment, which can be shared among several agents, with and through which it interacts. In conclusion, the environment of an agent can be seen as the union of its physical environment (the information it knows, or believes to be true, about its manufacturing and logistical environment) and its social environment (the information it knows, or believes to be true, from other agents).

Agent-based models, simulation, and systems

Having defined the concepts of agent, environment, and distributed systems, it is now possible to present the definitions of *agent-based models, simulation*, and *systems* that will be employed in this thesis.

An agent-based model of a real-world distributed system is built using the agent as the main modelling unit, and consists of three basic elements: a collection of agents, with their attributes and behaviours; some mechanisms of interaction between agents; an environment [40]. Also non-distributed systems (e.g. an automated production line) can be modelled with agent-based models, but this modelling technique fits better with distributed systems: the actors of these types of systems have in fact the fundamental characteristics of agents (i.e. situated, reactive, autonomous, etc.), which are, therefore, their perfect modelling twins. This is why the definition of agent-based models explicitly refers to distributed systems. It should also be noted that the term “system” appearing in the definition does not only refer to tangible real systems, but also to management systems. In this thesis both the mining system and the associated management will be considered, and two agent-based models will be presented: the first one (Chapter 5) is related to the as-is logistic and production system of the industrial partner (an underground gold mine); the second (Chapter 6) is the model of the system devised by the author as a solution to the fleet management problem.

Moving on to the next concept, the meaning of agent-based simulation can be defined by studying the etymology of the term. Simulation is a set of principles and techniques allowing to reproduce the behaviour of a real system, in order to respond to a specific problem of decision or design. The adjective "agent-based" denotes the fact of using agent-based models of the real system. As a result, in agent-based simulation the real system is first mapped into its respective agent-based model, which is then implemented on a computer system thanks to simulation platforms and/or programming languages, giving rise to a digital replica of the real system capable of virtually reproducing its behaviour. Considering what has been said for agent-based models, it is easy to understand that agent-based simulation is especially suitable for complex systems operating in dynamic environments: in contrast to traditional top-down approaches, the emphasis is in fact on the individual and on the interactions between individuals [39]. In this research project both the as-is mining system of the industrial partner and the to-be mining system designed by the author will be simulated using agent-based simulation.

Net of these first two definitions, the main benefit offered by agent-based technology comes out: it allows to *"model and simulate the dynamics of complex systems where several (decision-making) entities interact either directly or indirectly"* [41].

While the two concepts just revised refer to the modelling and simulation of real systems, agent-based systems deal with the support and automation of such systems. An agent-based system, in fact, is a computer system composed of a multitude of interacting elementary software units, whose individual characteristics are those of an agent. In fact, *"the fundamental characteristic of software agents is their ability to sense (i.e., have an internal perception of) their environment and to react (i.e., perform actions) in an autonomous manner (i.e., without the intervention of a human user) to changes in their environment"* [41]. Agents are therefore situated, reactive, and autonomous. In more advanced multi-agent systems, software agents can also be proactive and intelligent. Another key aspect characterizing multi-agent systems concerns the ability of agents to communicate with each other. A real-world example of these systems is provided by the Kiva robots (manufactured by Amazon Robotics) used in Amazon's warehouses, in which there is a software embedded in each vehicle capable of perceiving the environment and reacting, endowed with autonomy and behaviour, and having all the other features of an agent. In this research project, a tailor-made agent-based system will be the solution proposed by the author to the fleet management problem.

Having clarified the meaning of the individual concepts, we can now analyse how in this thesis they support one another to build an integrated whole. The integrated whole is represented in Figure 3.5, which outlines the main steps of the research project. In the first phase (*analysis*), the as-is mining system of the industrial partner will be investigated with agent-based modelling and simulation. Using, among other things, the results of the analysis, an agent-based system tailor-made for fleet management in underground mines will be conceived in the *design* phase. By implementing this solution in the starting (as-is) mining system, the to-be mining system is brought to life. This resulting cyber-physical system, however, will not be physically built in the real world, but will first be assessed (and this research project stops here) in the *evaluation* phase, where its agent-based model will be created and simulated with agent-based simulation.

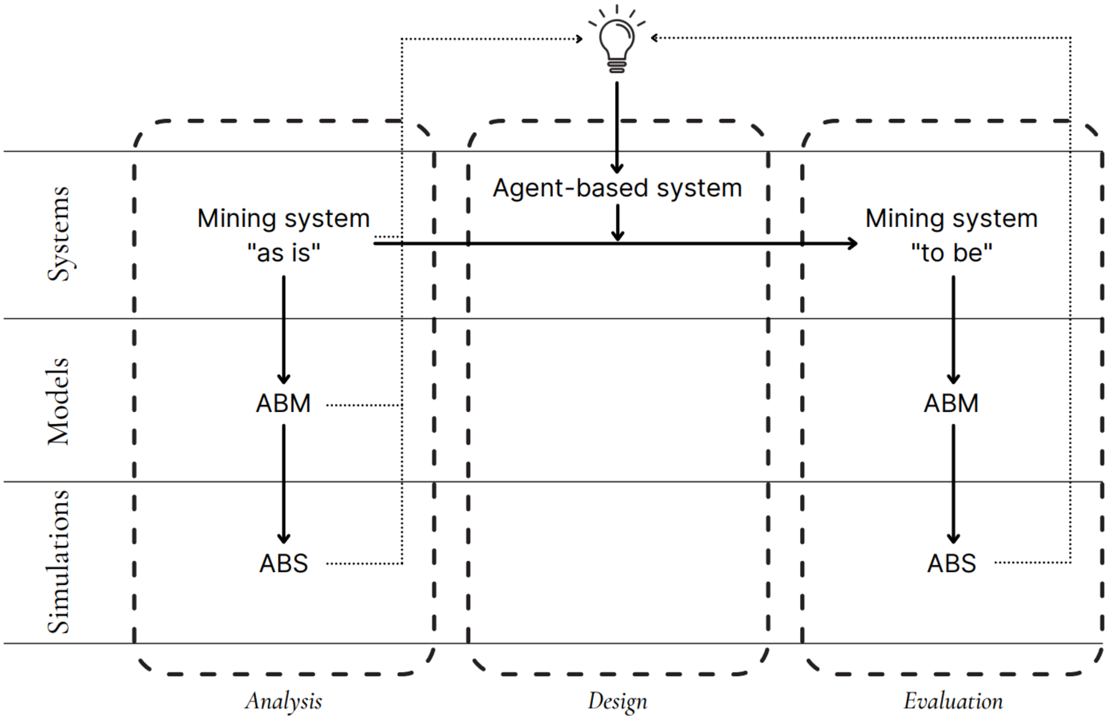

Figure 3.5 Deployment of agent-based modelling (ABM), agent-based simulation (ABS) and agent-based systems in the research project

CHAPTER 4 METHODOLOGY

4.1 Introduction

This chapter is structured as follows. First, emphasis will be placed on the type of research. The process followed to conduct the research will then be presented in Section 4.3. Finally, methodological choices and methods of data collection and analysis will be discussed in sections 4.4 and 4.5 respectively.

4.2 Type of research

When conducting research it is important to take into account the difference between fundamental and applied research [42]. Applied research is oriented towards solving practical problems, with motivations of pragmatic and utilitarian nature. Fundamental research, on the other hand, focuses on theory building, with the main motivation being to expand human knowledge. It is therefore straightforward to draw the link between applied research and engineering, and that between fundamental research and science. In the light of the foregoing, it is immediate to classify this thesis as applied, engineering research. All engineering research, moreover, is driven by the anticipated value of a future application. In the case of this research project, the term “application” refers to agent-based systems in the context of underground mines.

The question now is: how can we conduct engineering research? To answer this question, the author took as reference the distinction between science and engineering proposed by [3].

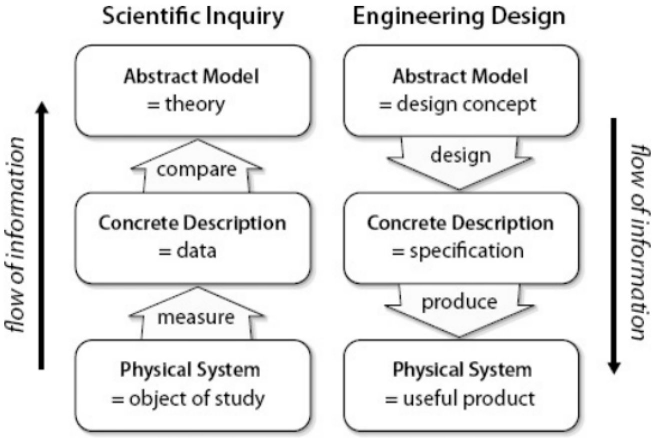

Figure 4.1 Scientific inquiry vs Engineering design by [3]

According to [3], the essence of science is inquiry, whereas the essence of engineering is design. The author distinguishes Scientific Inquiry and Engineering Design using the flow of information as a differentiator. Scientific Inquiry has a bottom-up structure: from studies and observations of the physical world, general theories are concluded; information flows from matter to mind. Engineering Design, on the other hand, has a top-down structure: from general knowledge and theories, engineers design specific solutions; information flows from mind to matter. In order to conduct engineering research, it seems rational to combine the bottom-up and the top-down structures. With regard to this thesis, more specifically, the baseline will be the knowledge about agent-based technology. The research project, in particular, will be triggered by the observation of how agent-based systems lend themselves very well to solving the fleet management problem in contexts other than underground mines, such as the industrial framework of AGVs. Starting from this knowledge, a multi-agent system tailor-made for fleet management in underground mines will be *designed*. The resulting cyber-physical system, however, will not be physically built (i.e. *produced*) in the real world, but will be implemented in the digital replica of an existing underground mine (industrial partner). Through simulation with AnyLogic software, it will then be possible to reproduce the behaviour of the conceived system and *measure* its performance. By using the inductive method, finally, the value proposition offered by the designed solution will be generalised to the entire world of underground mining. The initial theory on the capacity of the agent-based paradigm to manage the fleet in an underground mine will thus be confirmed, and new knowledge will be created. Bottom line, the flow of information will first go from the top to the bottom, and then back up again to close the circle.

4.3 Research design

Research design is defined by [43] as “the logic that links the data to be collected (and the conclusions to be drawn) to the initial questions of a study”. Thus, a good research approach must structure and connect research objectives, questions, data, activities, methods and results. The resulting, comprehensive research process followed in this thesis, in which all previous elements are meshed and support one another to form an integrated whole, is represented in the blueprint in Figure 4.2.

The research was initiated by a clarification phase, aimed at identifying the research focus. To do so, both the literature on underground mining and the needs of an industrial partner (a Canadian underground gold mine) were analysed. The mission of this thesis, in fact, is to create knowledge and at the same time propose practical solutions for industrial implementation. This dual perspective, i.e. industry-related and academia-related, was therefore

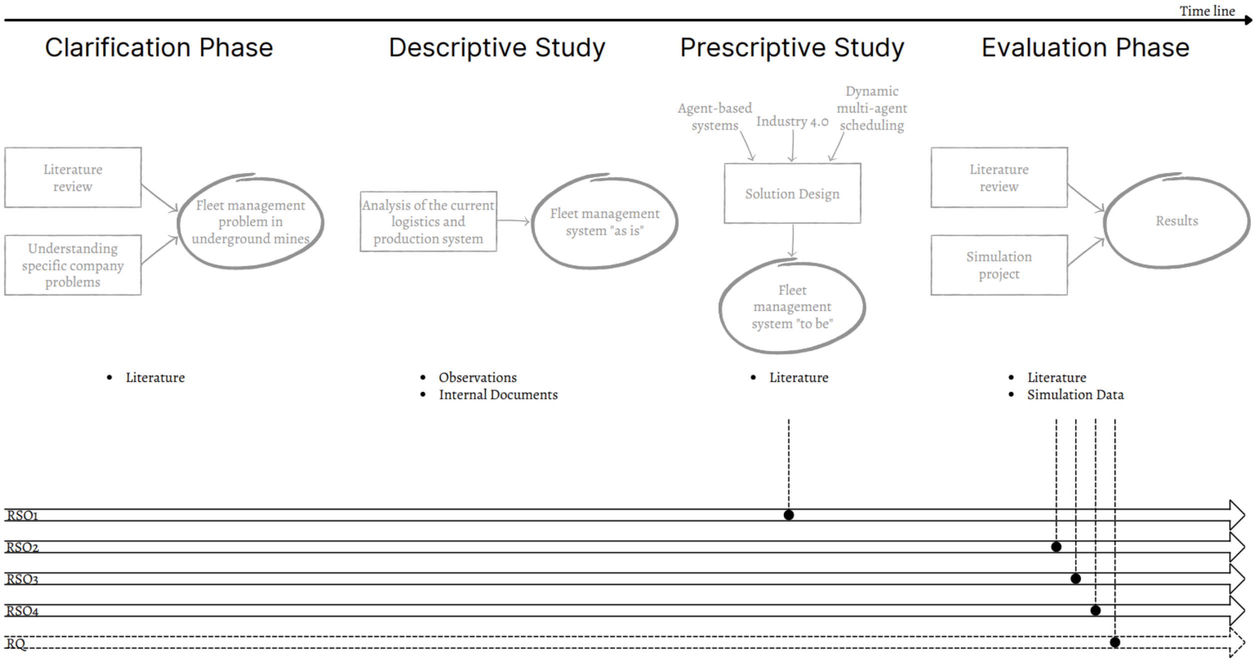

Figure 4.2 Research process

followed throughout the whole research project. The output of the first phase was the identification of the problem to be solved, i.e. the fleet management problem in underground mines. The latter is a well-known problem in the literature, in relation to which the industrial partner wishes to improve its performance.

In the descriptive study, then, an in-depth analysis of the current logistics and production system of the industrial partner was conducted. This shaped the *as-is fleet management system*, i.e. the benchmark of the *to-be fleet management system* that will be conceived in the next step (prescriptive study).

In the prescriptive study, the researcher was required to come up with innovative solutions to improve the present situation. Creativity thus played an important role at this stage. Knowledge of paradigms such as Industry 4.0, agent-based technology and dynamic multi-agent scheduling fed into the design of the solution, a multi-agent system tailor-made for the fleet management in underground mines. The literature on the employment of multi-agent systems to solve the fleet management problem, not only in underground mines, was also reviewed. No evidence of the use of such solutions in underground mining was found. Multi-agent systems, however, have already been applied to manage the fleet in industrial contexts such as that of AGVs, resulting in efficient and effective systems. This has encouraged the continuation of the research. In the next step of the process, therefore, the designed agent-based system (i.e. the fleet management system “to be”) was evaluated.

In the last phase of the research, the proposed solution to the fleet management problem was evaluated from both a quantitative and a qualitative point of view. As it is easy to guess, this phase was the most pivotal, as the success of the whole research depended on it. Simulation was used as an analysis and decision-support tool (it allowed the performance of the solution to the problem to be evaluated). By implementing the multi-agent fleet management system in the digital replica of the industrial partner mine, in fact, it was possible to assess the impact of the solution on the key performance indicators of the company. This made it possible to conduct a quantitative benchmarking between the to-be fleet management system (from the prescriptive study) and the as-is fleet management system (from the descriptive study). Last but not least, during this stage the multi-agent management system was also qualitatively evaluated, conducting a qualitative benchmarking between the proposed solution and: (i) the one currently adopted by the mine under review; (ii) the principal ones presented in the underground mining literature. An answer to the starting research question was finally found.

4.4 Methodological approaches

The roadmap of this thesis, as depicted in Figure 4.2, consisted of four main phases: a clarification phase which pinpointed the focus of the research, i.e. the fleet management problem in underground mines; a descriptive study that analysed the approach used by the industrial partner in relation to the research focus, shaping the as-is fleet management system; a prescriptive study in which the author designed an innovative solution for the investigated problem, shaping the to-be fleet management system (i.e. the multi-agent system); an evaluation phase of the proposed solution. While the key elements of the first three phases are ascertainment and creativity, the core of the fourth is analysis: in this last phase, therefore, several methodological choices had to be made. The rest of this section thus aims to explain which approaches and methodologies were adopted, and the rationale behind each of them.

The first decision that had to be taken concerned how to evaluate the performance of a solution (the designed multi-agent system) to a problem (the fleet management problem). To do so, several approaches were viable (Figure 4.3): one could have gone down the route of prototyping or experimenting with the real system, for example. Assessing the solution by creating a mathematical model, on the contrary, would have been utopian, due to the high complexity of the system under investigation. The approach chosen by this thesis was computer simulation: the multi-agent system was implemented in the digital replica of the real mine (built using AnyLogic simulation software), and simulation experiments were

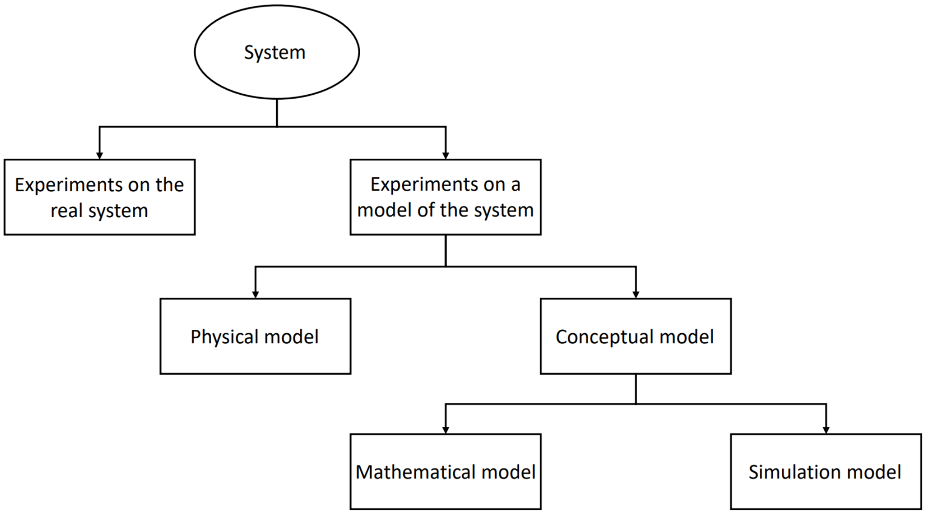

Figure 4.3 Different options for evaluating the performance of a solution to a problem

conducted to evaluate its performance. Simulation therefore enabled the goal to be achieved without having to interfere in any way with the real-world mine, and without having to invest money in prototypes of the solution.

Once simulation was chosen as a decision support tool, the next dilemma was the choice of which simulation technique to use. There are in fact four main techniques: Monte Carlo simulation; discrete event simulation; system dynamics; agent-based simulation. Agent-based simulation, in particular, was chosen for this research, as it is the paradigm most suited to the nature of the studied system. Mining systems are in fact distributed, as are almost all logistics and production systems. In a distributed system, there are: different actors in an environment over which they have a varying extent of control; decisions and actions of different nature made in parallel; interdependencies and interactions between actors, within and through the environment. All in all, the system behaviour is not centrally controlled, but emerges from the actions and interactions between these unitary entities. The most fitting paradigm for distributed systems is agent-based simulation: agents in agent-based simulation models are the perfect "digital twins" of actors in distributed systems (please refer to the discussion in Section 3.4.2). Evidence that agent-based modelling and simulation are the appropriate methodologies for this research project can also be found in the literature. [39], for example, argues that *"agent-based modeling is especially suitable for simulating the behavior of complex systems operating in dynamic environments"*, such as the mining system under consideration in this thesis. The individual behaviours, the interactions of individuals, and the emerging global (system-level) behaviour are in fact well captured with the agent-based paradigm, which is a decentralized, individual-centric (as opposed to system-level) approach [39].

The ensuing decision concerned the simulation software. AnyLogic [44] was selected, since: (i) the mining industry is one of the many sectors where this software has proven its modelling capabilities; (ii) it offers three simulation techniques (discrete event simulation; system dynamics; agent-based simulation), among which is the agent-based simulation sought by this research. The created simulation model, in this regard, was realised employing multi-method modelling, stemming from the amalgamation of agent-based modelling and discrete event modelling. As explained above, the agent-based paradigm is the cornerstone of this research, and it was the one predominantly used to realise the digital replica of the industrial partner. The discrete event paradigm, however, was the best suited to modelling some minor aspects of the logistics and production system: it was therefore employed, but with a much less pervasive impact on modelling.

Another parenthesis needs to be opened with regard to the methodology used to conduct the simulation project. To ensure that the study was undertaken with scientific rigour, and to eliminate the risk of the research ending up in unconnected streams, in which important steps in the project were skipped, the author used as a guideline the methodology proposed by [45]. This methodology consists of a step-wise approach comprising eight major phases one must follow for the proper application of simulation in industry. The eight stages are illustrated below. Each phase is further broken into steps (the reader interested in the comprehensive methodology and all the individual steps may refer to [45]).

Phase 1. Define the problem
Phase 2. Design the study
Phase 3. Design the conceptual model
Phase 4. Formulate inputs, assumptions, and process definition
Phase 5. Build, verify, and validate the simulation model
Phase 6. Experiment with the model and look for opportunities for design of experiments
Phase 7. Document and present the results
Phase 8. Define the model life cycle

The author opted to use this methodology as a framework since: (i) it is a well-known practice in academia and industry; (ii) it is a well-rounded methodology, able to embrace the whole life-cycle of a simulation project, starting from the up-front work to carry out prior to building the simulation model, all the way through to the management of the lifecycle of the created model; (iii) in addition to the factors traditionally covered by other methodologies (such as model coding, input and output data analysis, verification and validation), this methodology emphasises the roles of teamwork, project management, and collaboration with stakeholders in the simulation project, all of which are key aspects when conducting research in collaboration with an industrial partner, as is the case in this thesis.

Delving deeper into verification and validation of the simulation model, it is worth mentioning the associated methodologies employed in this thesis. Unfortunately, there is no standard procedure to follow, and there is no reference guide to determine which techniques to use each time: each simulation project presents a new and unique challenge to the model development team. Over the years, however, several approaches have been proposed for validation and verification in simulation studies. Most of them elaborate and boil down to the concepts propounded for the first time by [4]. For this reason the paradigm proposed by [4], presented in Figure 5.13 in Section 5.4.2, was adopted in this thesis. From a more pragmatic point of view, then, the practical approaches to verification and validation employed in this work refer to those proposed by [46]. The author decided to use [46] as a framework because this paper offers a broad and diversified portfolio of techniques and approaches, both qualitative and quantitative, that can be used for both observable and non-observable systems (where "observable" refers to the feasibility of collecting data on system behaviour). For a better comprehension of the meaning of verification and validation, and to appreciate how the paradigm proposed by [4] and the techniques and approaches proposed by [46] have been employed in this research, the reader is invited to consult Section 5.4.2.

4.5 Data collection & analysis methods

As already depicted with the blueprint in Figure 4.2, the data collection methods employed in this thesis are:

Literature review: Literature review was a central aspect of this thesis, as it served three purposes. First of all, an initial review of the state of the art in the management of logistics and production systems in underground mines was performed to clarify and better define the research focus (i.e. the problem to be addressed). Secondly, it enabled the author to find research gaps and shape improved research questions. Last but not least, literature reviews were used as benchmarks for comparing the proposed fleet management system with solutions proposed by other researchers.

Internal documents review: Production plans, dispatching lists for the working day under review (assigning to each vehicle in the fleet a list of destinations to visit), the layout of the underground mine, and documents on the vehicle fleet were the main company documents consulted.

Observations: This data collection method was central for conducting the analysis of the logistics and production system of the industrial partner and creating its digital replica using

AnyLogic simulation software.

Simulation experiments: The simulation experiments conducted with AnyLogic (presented later in the thesis) allowed to assess the impact of the proposed solution on the key performance indicators of the mine. The quantitative data generated enabled some of the main research objectives to be met, and to answer the research question underpinning this thesis.

With regard to data analysis methods, on the other hand, the author considers it appropriate to highlight the importance of making multiple replications in simulation experiments, which justifies the use of this practice throughout the rest of the thesis. By way of example, let us consider a random variable of interest, such as the utilisation of a specific truck in the fleet. Each replication of the simulation experiment will correspond to one observation of the random variable. The usefulness of making n replications, therefore, consists in: (i) obtaining n observations of a random variable of interest; (ii) ensuring that the n observations are independent and identically distributed. The latter is a hypothesis very demanded in statistics, for example when one wants to compute confidence intervals, as in the case of this thesis. The determination of confidence intervals, moreover, requires knowing the nature of the random variable, i.e. its probability law. To make data react as if it came from a normal distribution, the Central Limit Theorem was invoked.

CHAPTER 5 SIMULATION MODEL OF A REAL UNDERGROUND MINE

5.1 Introduction

This chapter will present the logistics and production system and the related management approach of the industrial partner (an underground mine). The dichotomy mining system-management system will then be mapped into a simulation model.

Delving into details, this research has been conducted in collaboration with a Canadian underground gold mine. During the early stages of the project, the mining system of the company was observed and analysed. Particular emphasis was then placed on the decisions that the firm takes in order to manage the aforementioned logistics and production system. Finally, the author entangled the analysed mining system and the adopted management policy in an agent-based simulation model of the real mine, giving life to the digital replica of the industrial partner. These three steps of the research project will be outlined in detail in the following sections.

5.2 Industrial partner: mining logistics and production system

The purpose of this section is to present the logistics and production system of the mine under study. The latter consists, in a nutshell, of a heterogeneous fleet of vehicles that move between specific points in the mine and perform specific operations at each point. First of all, the environment of the mining system, i.e. the transportation network, will be analysed. Subsequently, the main actors of the logistic and production system, i.e. the fleet vehicles, will be defined. Finally, the functioning of the mining system will be presented using two standpoints, a resource-based perspective and an activity-based perspective.

The transportation network

The nature of the haulage network of an underground mine, as already anticipated in the previous chapters, is a major constraint of the fleet management problem. It consists, in fact, of one-lane bi-directional road segments, which constrain the movement of vehicles. In order to better understand the meaning of these terms, the layout of the mine under investigation will be analysed below. It should be noted that only parts of the entire underground transportation network will be considered during the simulation project. In this thesis, in fact, only a sub-set of the entire vehicle fleet, consisting of around a hundred vehicles in total, will be taken into account: only the branches of the transportation network falling within the competence of the selected fleet will therefore be considered and modelled in the digital model of the mine.

In an underground mine there are several underground levels "parallel" to the Earth's surface, each at a different depth, where the main activities take place. It is possible to reach these levels by means of a ramp, which connects each of them to the surface. Both the ramp and the levels are made up of tunnels that can accommodate only one vehicle in width, and in which vehicles travel in both directions. Hence the term one-lane bi-directional road segments. Occasional passing bays, where a vehicle can temporarily stop to let the oncoming vehicle pass, play therefore a crucial role. Going into further detail, ten levels of the mine under review have been modelled, as they are within the purview of the considered vehicle fleet. In contrast to the underground network, the ground floor, i.e. the one on the Earth's surface, does not feature any transportation network restricting the movement of vehicles. By means of the ramp, trucks loaded with ore can reach the surface, where they can freely move to reach unloading points and feed the above-ground plant, where the first processing activities of the raw material can thus begin. The investigated mine also features a second ramp, which flows into the first one and which is not directly connected to any level. This ramp is used by trucks with maximum transport capacity, and leads them to a second plant on the surface.

The fleet of vehicles

Having defined the environment of the logistics and production system, i.e. the transportation network, it is now appropriate to focus on who the main actors of the system are, i.e. the vehicles of the fleet. The overall fleet of the company consists of many vehicle families, each with a specific function and homogeneous within itself. The total number of vehicles is considerable, but for the purposes of this research only part of the entire fleet will be addressed: considering the research objectives, in fact, the behaviour of this simplified mining system has been estimated to be a good proxy for the behaviour of the real system including the entire fleet.

Table 5.1 shows the vehicle families considered, specifying the number of units in each family, the activity performed, and the mining process in which each vehicle type is involved (for a better understanding the reader may refer to Section 1.2). The behaviour of each type of vehicle and the description of each activity presented in Table 5.1 will be described in the next section, which will deal with the functioning of the logistics and production system of the underground mine.

Table 5.1 The modelled fleet of vehicles

Vehicle family	Number of units	Activity	Process
Drillers	3	Drilling	Development process
Chargers	3	Charging	
LHDs (Development)	2	Loading	
Trucks (Development)	4	Ore transportation	
Bolters	6	Bolting	
Cleaners	1	Cleaning	
LHDs (Production)	3	Loading	Production process
Trucks (Production)	6	Ore transportation	
Personnel Carriers	1	Pick up workers	End-of-shift operations

Functioning of the mining logistics and production system

Having defined the actors and the environment of the mining system, we now have all the tools needed to understand its functioning. In order to present as clear a picture as possible, two points of view will be employed: a resource-based perspective and an activity-based perspective.

Resource-based perspective

With this first perspective, the behaviour of the mining system will be studied by taking the point of view of its resources, i.e. the fleet vehicles. The duty of the different mining vehicles is to move towards selected points in the mine, and, once reached the destination, to perform some operations. More specifically, a first group of vehicles consisting of drillers, chargers, bolters and cleaners present a simple logic of functioning. After having completed the current mission, i.e. after having reached an established underground point and having carried out the operation they are responsible for, they are ready for a new mission towards a new underground destination. What distinguishes the vehicles in this group is simply the nature of the performed activity. Trucks and load-haul-dump machines (LHDs), on the other hand, exhibit a more articulated behaviour and an increased need for coordination. Once the underground destination has been reached, in fact, the loading activity can only start once both vehicle types are present. Once loaded, the truck then starts its trip towards one of the two above-ground industrial plants, depending on its transport capacity: in order to achieve transport economies, trucks with a higher transport capacity are dispatched to the second plant, which is further away from the underground working areas than the first plant. Once it reaches the surface, the truck moves towards certain points in the mill (waste dumps and stockpiles), according to the nature of the cargo being transported (waste, low grade ore, high grade ore, super-high grade ore). After being unloaded, the truck can then return to its underground starting point, to be loaded again by the LHD, and begin a new journey to the surface. This cycle continues until the entire bulk in the underground work area has been removed. Only at this point the LHD and the truck can start a new mission towards a new underground destination.

Activity-based perspective
The behaviour of its resources having been highlighted, we can now study the mining system from another point of view, that of the underground working areas. As already explained in the transportation network description, the mine consists of several underground levels, each with a different depth, accessible by means of two ramps; each level, more specifically, contains several sites, which require certain activities to be performed by the fleet of vehicles. Underground sites are therefore the "actors" that played the role of vehicle destinations in the resource-based perspective. Two types of mining processes take place in an underground mine: the development process and the production process. Development mining means excavating in order to gain access to the orebody, while production mining means excavating and extracting ore from the orebody which has been accessed and made available by development. A given site at a given time, therefore, engages in one of two processes. As already anticipated, this thesis will focus on the entire development cycle, while as far as production is concerned, only the loading and transport activities performed by LHDs and trucks will be modelled. Hereafter, therefore, we will put ourselves in the shoes of a site dedicated to the development process. Advancing the development process in a site means iterating the so-called development cycle in that work area. The structure of the development cycle has already been illustrated in Section 1.2 of this thesis (Figure 1.1). It consists of an ordered sequence of activities, each carried out by the corresponding vehicle. The objective of each cycle is to increase the length of the underground tunnel. Drillers carry out the drilling activity, in which the rock face is drilled in order to create holes, which will then be filled with explosives by the chargers. Chargers are therefore responsible for loading the explosive. Once the detonation has taken place, the exploded bulk is removed from the work area by the LHD and trucks. Bolters then install protective nets on the walls of the newly exploded tunnel to ensure safe operations in that section of the mine. Ultimately, cleaners are in charge of removing debris from previous operations and making the work area ready for the next development cycle, which will start all over again with the drilling activity.

5.3 Industrial partner: management system

In the previous section a picture of the logistics and production system of the gold mine under scrutiny was presented. The objective is now to analyse the system adopted by the company to manage such a mining system. This section, more specifically, will present the tools, the approaches and the overall strategy implemented by the firm for planning, scheduling, execution and control. More precisely, with reference to the different management levels of a logistics and production system (in Figure 5.1), this thesis will focus only on operational planning, scheduling, and execution and control, in line with the research objectives.

Figure 5.1 Planning levels of a logistics and production system

With regard to operational planning, the production plans generated by the company will be illustrated. As regards scheduling[1], the dispatching decision-making approach will be investigated (please remember: the dispatching decision concerns the choice of the next destination for a vehicle that has just finished its current task). As far as execution and control are concerned, we will analyse how the company and its actors take routing and scheduling decisions (please remember: the routing decision refers to the choice of the path to be taken to reach the selected destination, while the scheduling decision concerns the management of traffic in the mine).

Operational planning - Production plans

With a rolling approach of weekly frequency, the mine conceives two main production plans, related to the development and production processes respectively. The main inputs to this operational planning phase are the strategic plans, the tactical plans, and the production

---

[1] Note: The term "scheduling" is used in this thesis with different meanings depending on the context. When the emphasis is on Operations Management, the term refers to the scheduling of operations. When the spotlight is on underground mines and the fleet management problem, instead, the term refers to traffic management in the underground transportation network, in accordance with the literature.

performance of the mine in the week before the one to be planned. With regard to production planning related to the development process, Figure 5.2 shows a mock-up of a real production plan generated by the mine.

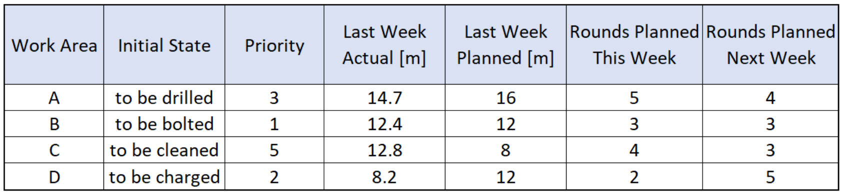

Work Area	Initial State	Priority	Last Week Actual [m]	Last Week Planned [m]	Rounds Planned This Week	Rounds Planned Next Week
A	to be drilled	3	14.7	16	5	4
B	to be bolted	1	12.4	12	3	3
C	to be cleaned	5	12.8	8	4	3
D	to be charged	2	8.2	12	2	5

Figure 5.2 Development plan: a simplified template

As we can note, production targets are set with a planning horizon of two weeks, and the planning process is performed adopting a rolling approach of weekly frequency. The production plan establishes which sites will be subject to the development process during the planning horizon, and how many development cycles will have to be performed for each of them. The plan features a high level of aggregation, as the number of rounds to be performed is set for the whole week, without detailing the daily objectives. A lower level of aggregation, on the contrary, can be found in the production plan related to the production process (Figure 5.3 depicts a mock-up), which specifies the objectives to be achieved on a daily basis.

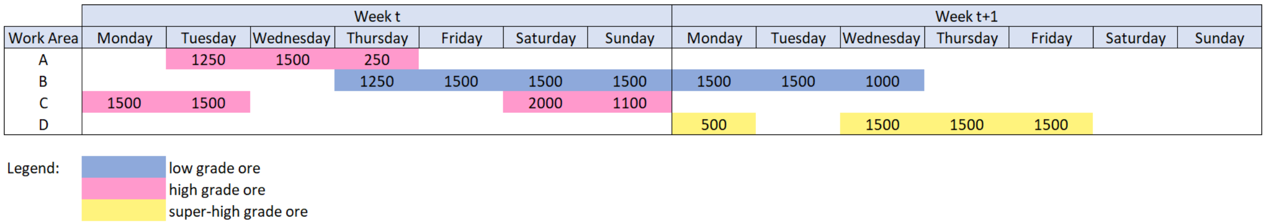

	Week t							Week t+1						
Work Area	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday
A		🟪 1250	🟪 1500	🟪 250										
B				🟦 1250	🟦 1500	🟦 1500	🟦 1500	🟦 1500	🟦 1500	🟦 1000				
C	🟪 1500	🟪 1500				🟪 2000	🟪 1100							
D								🟨 500		🟨 1500	🟨 1500	🟨 1500		

Legend:
🟦 low grade ore
🟪 high grade ore
🟨 super-high grade ore

Figure 5.3 Production plan: a simplified template

The production plan in Figure 5.3 defines which draw points will be subject to the production process over the planning horizon, specifying the quantity to be "produced" at each of them, on a daily basis. Information on the nature of the ore produced at each point is also provided. The gold resulting from the production process, in fact, is characterised by a certain degree of purity: the mine classifies the ore in low, high and super-high grades.

An in-depth analysis of the two production plans just presented was essential to the development of this research project, as they are the main inputs to the fleet management problem. In other words, the dispatching, routing and scheduling decisions addressed by this thesis cannot be taken without considering the production targets defined during the operational planning process.

Scheduling - Dispatching decision-making approach

In the following, the author will illustrate how the mine under investigation takes dispatching decisions, which consist of defining the next destination for a vehicle that has just completed its task. At the beginning of each working day (consisting of two 12-hour shifts), engineers assign to each vehicle in the fleet a list of destinations for the entire working day, taking into account the targets set for the week by the production plans. Each mining vehicle, therefore, once finished with its current operation at a site, is already aware of the next destination: it will look at the list of destinations that was provided at the beginning of the shift, and will head for the destination immediately after the one just treated. The dispatching decision of the fleet management problem, therefore, is not taken dynamically, but by solving a major initial dispatching problem. In other words, every time a vehicle finishes an activity, the dispatching decision is a sunk decision already made in the past, and is not dynamically evaluated.

This management approach is not ideal in the world of underground mining, due to the high uncertainty in which the mining logistics and production systems live. Having a list of tasks to be accomplished throughout the whole working day, in fact, is not in line with the requirements arising from living in a highly uncertain world. Evidences of the above can also be found in [6,14,15]. In the course of this thesis, therefore, the dispatching decision-making approach will be the main element of the management system that the author will seek to improve, relying on agent-based paradigms and dynamic multi-agent scheduling.

Execution and Control - Routing and scheduling decision-making approaches

During the execution phase of logistics and production activities, routing and scheduling decisions have to be made. Given a pair of points (a departure point and an arrival point) in the transportation network, and given alternative routes leading from the departure point to the arrival point, making the routing decision means choosing which path to follow. Scheduling decisions, on the other hand, have to be made because of the constraints imposed by the transportation network. The one-lane bi-directional road segments, in fact, constrain the movement of resources within the underground mine: while travelling, vehicles must constantly make decisions to coordinate with others, in order to avoid traffic conflicts (e.g. head-on collisions).

The company delegates routing and scheduling decisions to vehicle drivers. As far as routing is concerned, this means that, when several alternative paths exist to reach a given point, it is up to the driver to choose which one to follow. As far as scheduling is concerned, instead, decisions are delegated to workers with a lower degree of freedom. Whenever two vehicles with opposite directions are about to collide (head-on conflict), in fact, drivers are free to coordinate as they see fit to resolve the conflict, but, where possible, they should adhere to the policy of giving priority to loaded trucks. Human decisions, together with this rule, will define which of the two vehicles in the head-on conflict will have to back up and temporarily stop in the passing bay to let the oncoming vehicle pass. A clearer picture of how mine actors make routing and scheduling decisions will be presented in Section 5.4.1. What is important to understand for now is that these two decisions are outputs of the individual worker's decision-making process. Section 5.4.1, in this regard, will present how the author decided to model such decisions made by humans. In the improved management system proposed in this thesis, to conclude, routing and scheduling decisions will no longer be delegated to drivers, but to software that will control each vehicle (the so-called software agents).

5.4 Industrial partner: digital replica

5.4.1 Agent-based simulation model

The creation of a simulation model of the mine under investigation is a milestone in this research project, as one of the key research objectives is to employ simulation to reproduce the behaviour and evaluate the performance of the proposed fleet management system. In the previous two sections, a thorough description of the logistics and production system and of the associated management system has been provided. The union and integration of these two systems shape the part of the company whose digital replica is going to be created. To achieve this goal, first an agent-based model of the real mine will be designed, then the latter will be translated into a simulation model implemented in AnyLogic. At that point we will possess a digital model of the real mine, capable of reproducing its behaviour by running the simulation, and in which risk-free experiments can be conducted. This section, more specifically, will deal with the first and main step of the path towards the creation of the digital copy of the mine, namely the modelling phase. In the following, therefore, the reader will be presented with the agent-based model of the mine envisioned by the author.

Table 5.2 Endogenous and exogenous elements of the problem

Endogenous elements	Exogenous elements
Dispatching decision	**Production plans**
Routing decision	**Vehicle fleet**
Scheduling decision	**Transportation network**
Maintenance strategy	**Development process**
Fuelling strategy	Production process
...	...

Before conceiving the agent-based model, a propaedeutic study was conducted: the problem addressed in this thesis was analysed with the purpose of identifying the needs for the modelling phase. The output of this analysis was the identification of the endogenous and exogenous elements of the problem (the most significant ones are presented in Table 5.2). The items in bold in Table 5.2 are the elements of the problem that will be modelled, while the others will be neglected, since outside the focus of the research. The main element to be improved with this thesis is the fleet management system. Managing a fleet of vehicles in an underground mine, more specifically, means making decisions of a different nature, including operational decisions (e.g. deciding destinations, routes, and traffic coordination for vehicles), maintenance decisions, financial decisions, etc. In this work we will only address the dispatching, routing and scheduling decision-making approaches, without calling into question the strategies of maintenance, refuelling, etc. The list of exogenous elements in bold, on the other hand, defines the problem components that will be modelled, but over which the author will have no control. Although modelled, in fact, the production plans, the fleet (e.g. number of vehicles), the transportation network (e.g. location of passing bays), and the development processes (sequences of activities to be performed) will not be questioned. Thanks to this initial analysis, the needs for the modelling phase have been pinpointed. The agent-based model of the mine can therefore begin to be designed.

As already mentioned in the previous chapters, agents are the backbone of an agent-based model (please see Section 3.4). In the model conceived by the author, each mining vehicle has been modelled as an agent, i.e. as an instance of a given agent type (just as each of us can be considered an instance of the “person” agent type). By definition, an agent is located in an environment from which it perceives information and on which it acts. The environment of each vehicle agent, therefore, includes all the other vehicles in the mine, the transportation network in which it moves, the dispatching list assigned by the mine engineers, and all other perceived information. The remainder of this section will outline: (i) the main agent types in the model, specifying roles, responsibilities, attributes, states, behaviours, and interactions with other agents; (ii) agent behaviour in relation to dispatching, routing and scheduling decisions, on which particular emphasis will be placed as they are the focus of this research project.

Vehicle agent type (superclass)

All agent types in the model, described in the following paragraphs, are subclasses of the *vehicle* agent type, i.e. they inherit from the latter the behaviour, the attributes and the interaction mechanisms. To better understand the meaning of inheritance, the following analogy can be taken as a reference. If we consider traffic in a city, although there are different types of vehicles (e.g. cars, buses, trams), each of them has the same driving behaviour (e.g. driving in the right lane), the same attributes (e.g. speed), and the same interaction mechanisms (e.g. using the blinker before turning). Although buses, for example, have specific functions (e.g. picking up passengers at stops), specific behaviours (e.g. waiting a certain time at stops), specific attributes (e.g. maximum capacity), and specific interaction mechanisms (e.g. validating passenger tickets), all these specific elements are in addition to and extend the aforementioned features common to all vehicle types. The exact same situation arises in an underground mine: each vehicle type inherits the behaviours, attributes and interactions defined in the superclass *vehicle*, where the characteristics shared by all vehicle types are defined, and extends the latters with its own specificities.

Examples of attributes of the *vehicle* agent type, common to each mining vehicle, are the speed and the priority index in traffic (loaded trucks have priority over empty trucks and all other types of vehicles, as happens in the real mine).

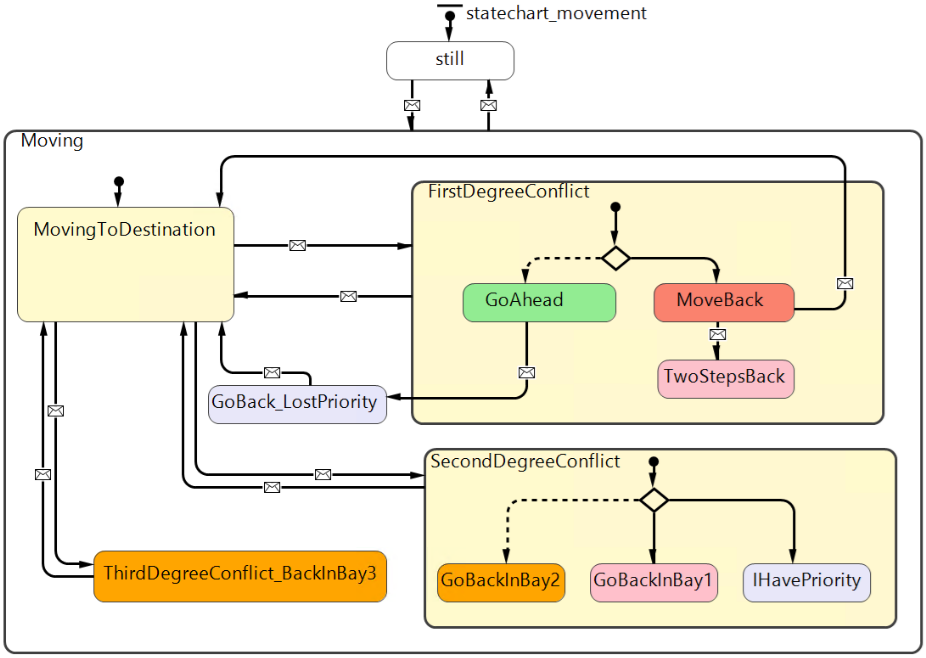

Figure 5.4 Statechart of the *vehicle* agent (implemented in AnyLogic)

The statechart in Figure 5.4 was implemented in AnyLogic and defines the states and the behaviour of a *vehicle* agent. Such an agent, in particular, may be either still (e.g. when working at an extraction point) or moving in the transportation network. While moving towards a destination, the road may be clear (the agent is in the MovingToDestination state) or there may be traffic conflicts with other vehicles of the fleet (the agent is in one of the other inner states of the Moving state). The presented statechart is able to handle up to 4 vehicles in conflict on the same section of the haulage network (an example of such a situation is provided in Figure 5.5).

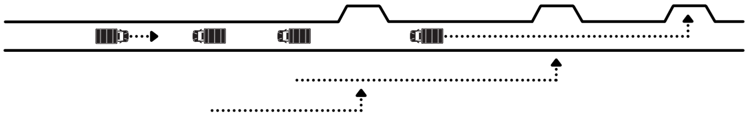

Figure 5.5 Example of a conflict involving 4 vehicles

The specific scenarios and the associated states will not be further discussed, but the overarching logic underlying the statechart will be explained below. The general idea is to enter into conflict as soon as a moving vehicle perceives a vehicle coming from the opposite direction. At this point it must be decided which of the two vehicles will take precedence over the other, which will thus reverse to the nearest passing bay, let the other pass, and finally resume its journey. In order to understand how the agents of the simulation model make this decision, reflecting the decision-making processes of the real mine drivers, the reader can refer to the end of this chapter, in which the agent behaviour related to the scheduling part of the fleet management problem will be addressed. In any case, once the “winner” of the conflict is chosen, the counterpart starts moving back to the first available passing bay (MoveBack state), where it will wait for the oncoming vehicle (in the GoAhead state) to pass. At this point the conflict will be resolved, and both agents will re-enter the MovingToDestination state, resuming their natural course to the destination.

Equipment agent type

This agent type is a subclass of the *vehicle* agent type. Drillers, chargers, bolters and cleaners are all modelled as *equipment* agents, as they all play the same role of reaching an underground site and performing the corresponding operation (drilling, charging, bolting, cleaning). Each agent is provided with a list of destinations to visit during the work shift (assigned by the engineers at the beginning of the working day). In addition to the behaviour inherited from the superclass, these agents exhibit the following functioning: as soon as the worker turns on the vehicle at the beginning of the shift, the agent starts the first trip planned in the work list; once arrived at destination, it may happen that the site is still being worked by the previous vehicle in the development cycle (e.g. a charger arrives at destination but the drilling activity is not finished yet): in this case, the agent will wait for the end of the ongoing operation; after the potential waiting period, the agent will start its operation, and, once finished, will start the working cycle again by moving to the next destination on the list. The statechart implemented in AnyLogic that refers to this behaviour is presented in Appendix A.

In addition to the traffic interactions inherited from the *vehicle* agent type, *equipment* agents interact with the *carrier* agent. At the end of the work shift, in this regard, the personnel carrier is in charge of picking up the workers in the underground mine and transporting them to the surface. For the blast to take place, in fact, all workers must be outside the mine. At the end of the shift, therefore, the *equipment* agents interact with the *carrier* agent, who leads them to the shutdown by picking up the workers. An analogous situation occurs at the beginning of the shift: after the blasting, the personnel carrier drives the workers to the underground mine, leaving each of them in proximity of the corresponding vehicle, which is thus brought back into service.

LHD agent type

Both the load-haul-dump (LHD) loaders involved in the development process and those involved in the production process were modelled with this type of agent. LHDs are in charge of reaching an underground site, loading ore/waste from it, transporting the material to trucks placed near the site, and dumping it into the trucks. As already mentioned, this type of agent inherits from the *vehicle* agent type. It features, however, specific attributes such as the transportable bulk capacity, expressed in tonnes. *LHD* agents, like *equipment* agents, are equipped with a list of destinations to be visited during the day, and have the exact same interaction mechanism with the carrier. These agents, moreover, are governed by the same behaviour presented for *equipment* agents, the only exception being the need for coordination with trucks. After arriving at the site and waiting, if necessary, for the end of the operations still in progress, the LHD can perform the loading activity only if the truck is also there, ready to be loaded. Therefore, while in the case of *equipment* agents a single resource was able to perform the corresponding operation on its own, in the case of the loading activity the collaboration between two mining resources is required. The statechart implemented in AnyLogic is presented in Appendix A.

Truck agent type

Both the trucks employed for the development process and those employed for the production process belong to this class of agents. *Truck* agents are responsible for concretising the efforts made by all other agents, yielding the output of the underground mine: they transport the ore obtained during the underground process to the surface, and thus feed the plant, which can start processing the raw material. In addition to the attributes inherited from the *vehicle* agent type superclass, *truck* agents are characterised by properties such as the transportable bulk capacity, expressed in tonnes. Similarly, as an extension of the inherited driving behaviour, these types of agents exhibit the operating logic illustrated in Figure 5.6.

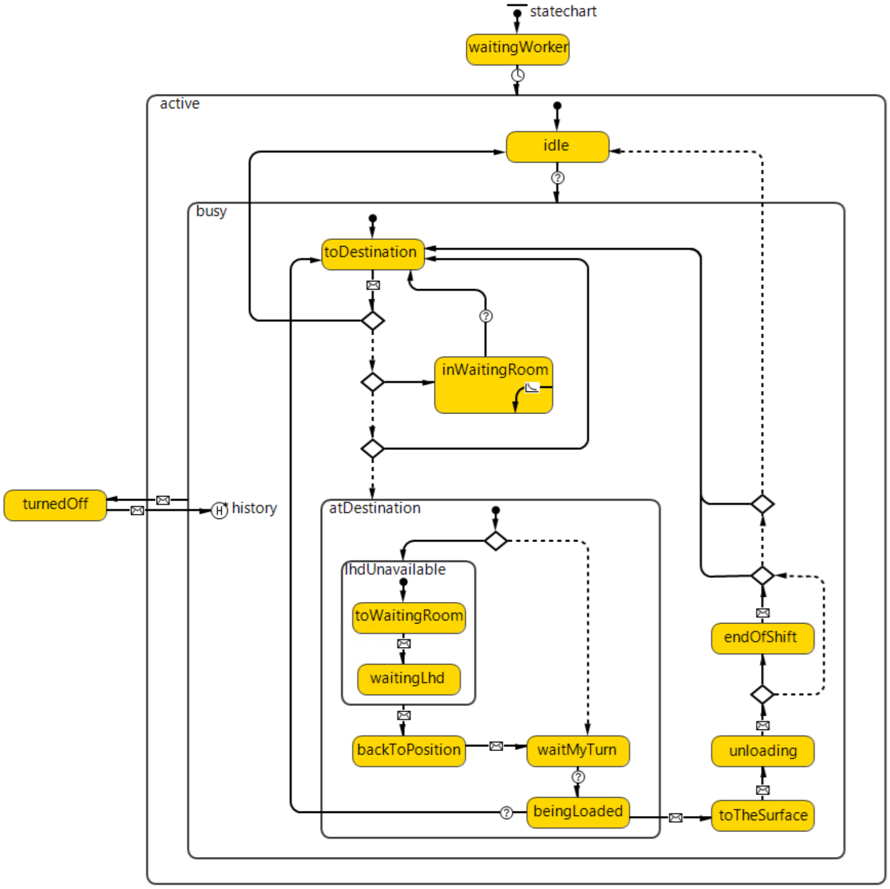

Figure 5.6 Statechart of the *truck* agent (implemented in AnyLogic)

The statechart in Figure 5.6 (implemented in AnyLogic) defines the event- and time-driven behaviour of the *truck* agent. After the journey towards the destination site defined by the work schedule provided by the engineers (toDestination state), it may be necessary to wait for the end of the previous activities (inWaitingRoom) and/or for the arrival of the LHD (lhdUnavailable state), since both types of vehicles must be in place for loading to take place. The mine, furthermore, adopts a fleet management policy whereby two trucks and one LHD are assigned to a loading operation, whence the presence of the state "waitMyTurn", in which the agent waits for the LHD to finish loading the other truck in place, if necessary (i.e. if also the other truck is in place at that time). Once loaded with ore or waste rock, the truck can start its journey to the above-ground plant (plant number 1 for trucks involved in the development process, plant number 2 for trucks involved in the production process). Once arrived at the surface, the agent, depending on the nature of the freight transported (waste, low-grade ore, high-grade ore, super-high-grade ore), heads for the corresponding discharge point. After unloading, if the end of the work shift is approaching, the truck remains on the surface waiting for the next shift to start; otherwise the truck will re-enter the underground mine and drive to the previously visited site, in order to be loaded again.

The *truck* agent interacts with the carrier, in a very similar way to that presented for *LHD* and *equipment* agents: once the *carrier* agent picks up the driver of the truck in question, the truck will be switched off (TurnedOff state), and will become operational again only with the arrival of the next shift worker. Underpinning the social nature of the *truck* agent are also interactions with the *LHD* agent. As soon as all the material at the site under processing has been picked up and removed from the underground work area, in fact, the *truck* agent will receive a message from the *LHD* agent informing it that the current task has been completed: the agent will thus start a new journey to the next destination on the dispatching list.

Carrier agent type

As already hinted, this agent becomes operational at the end of the work shift, and is in charge of picking up the workers in the underground mine and bringing them to the surface. Once the agent arrives in proximity of the single vehicle, the worker turns off the vehicle and gets on the carrier. After picking up all the workers, the agent returns to the surface and gives the approval for the detonation to take place. After blasting and ventilation of the toxic fumes, the carrier brings the workers back into the underground mine, and the new work shift can thus begin.

Agent behaviour: dispatching, routing and scheduling decisions

The agents introduced above need to make some decisions in certain situations, i.e. in a specific internal state and/or in a state of the environment: whenever a vehicle finishes working at a site, its next destination must be defined (dispatching decision); whenever it reaches a point in the mine from which several alternative routes can be taken, it must decide which path to follow (routing decision); whenever it is involved in a conflict with other vehicles moving in the network, it must coordinate with the others by deciding which one will step back to let the other pass (scheduling decision). In the following, the author will elaborate on the agents' behaviour in relation to these three decisions, since the latter are the main pillars of this thesis. *Equipment* agents, *LHD* agents and *truck* agents, in this respect, all exhibit the same decision-making processes, which have therefore been defined in the superclass *vehicle* and inherited by the three agent types.

As far as dispatching is concerned, the agents do not present a very sophisticated behaviour, as this decision is taken by the mine engineers at the beginning of the work shift. This means that whenever a vehicle finishes the current task, the next site/assignment is determined simply by reading the list of destinations assigned by the engineers at the beginning of the shift.

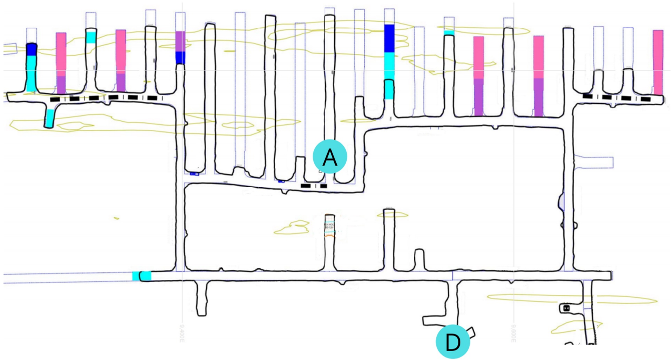

Figure 5.7 Routing decision: which route to follow to get from D to A?

With regard to routing, we can take the example illustrated in Figure 5.7 as a case study. The figure shows a situation in which a vehicle has two alternative routes available to reach point A (arrival) from point D (departure). Workers at the mine under examination stated that they made such a decision based on the length of the paths and on the state of the underground mine. This means that, under normal operating conditions, the worker will choose the shortest route. Due to the high uncertainty inherent in underground mines, however, it can happen that the longest routes are chosen by drivers: if we assume, for example, that there is a breakdown in a vehicle that is travelling the shortest route, all other cars will be shunted to the alternative routes. For the purposes of this research, the real mining system will be studied, at least from a quantitative point of view, under normal operating conditions; contingencies will in fact only be considered in the qualitative benchmarking between the real and the proposed mining system (see discussion on agility in Chapter 7.3). The digital replica of the real mine, therefore, can be seen as an ideal world in which *equipment* agents, *LHD* agents, *truck* agents, and the *carrier* agent will always choose the shortest of the alternative routes.

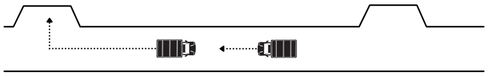

Figure 5.8 Scheduling decision: which vehicle in the head-on conflict should back off

While dispatching and routing decisions are governed by simple rules of behaviour, a much more articulated decision-making process is associated with scheduling. Let us consider the case in Figure 5.8, in which two vehicles are moving in opposite directions on the same stretch of road, and let us analyse how agents make decisions to resolve the conflict. It is worth noting that the created simulation model, in reality, is capable of handling conflicts involving up to 4 vehicles at the same time (an example is provided in Figure 5.5): the functioning logic, however, will only be provided for the simplest case involving two vehicles, since in all other cases the reasoning will be exactly the same, with the sole exception of having more stakeholders to coordinate.

Returning to the situation in Figure 5.8, the decision to be taken concerns the choice of the vehicle that will have to move back and stop in the bay to let the other vehicle through. Such a scheduling decision for the conflict resolution is delegated to individual workers, and the only rule of behaviour established by the company consists of giving priority to loaded trucks. This means that workers can coordinate as they see fit to resolve the conflict, obeying, when applicable, the aforementioned policy. Human decisions, therefore, are at the very heart of the functioning of the real world that the agent-based simulation model is seeking to simulate. This made it necessary in the modelling phase to devise a method for reproducing the human decision-making process. Figure 5.9 , in this regard, depicts the decision diagram with which the author modelled the decision process of the mine workers.

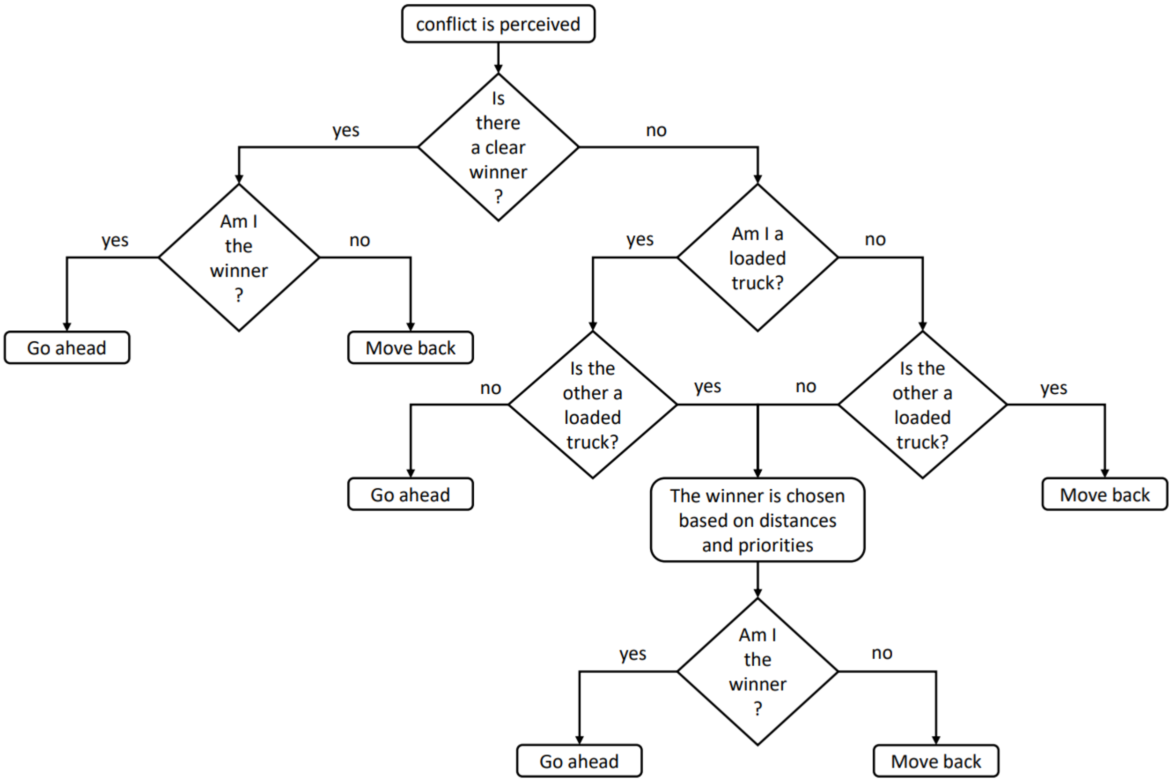

Figure 5.9 Decision diagram to model the decisions of real-world actors

The decision diagram represents the reasoning followed by the mine workers in deciding which of the two vehicles should back down to resolve the conflict in Figure 5.8. The first part of the diagram can actually be considered as a non-decision, but a simple reaction of the workers to the perceived environment: this is the case of conflicts where there is a clear “winner” (one vehicle is very close to the related bay while the other is very far away), and where the drivers, being aware of this, react accordingly. If, on the other hand, a solution to the conflict is not obvious, the workers first rely on the policy adopted by the mine: loaded trucks will have priority over empty ones and all other types of vehicles. If such a co-ordination rule is not able to manage the conflict (e.g. the two competitors are vehicles other than trucks), it is up to the workers to make a decision (lower part of the right-hand branch of the diagram). Workers at the surveyed mine stated that they decide on the basis of two factors: the distance to the nearest bay, to be travelled in reverse to resolve the conflict; the degree of priority of the mission in progress. Logically, the closer a vehicle is to the respective bay, and the lower the priority of its task, the more that vehicle becomes a candidate to move back to let the other one pass. Considering the two factors (obviously approximately), the drivers choose which of the two is the winner of the conflict.

The decision-making process of real-world actors was thus modelled with the decision diagram presented above. The latter was then translated into the behaviour of the agents of the simulation model (in Figure 5.10), that aim to imitate the aforementioned actors.

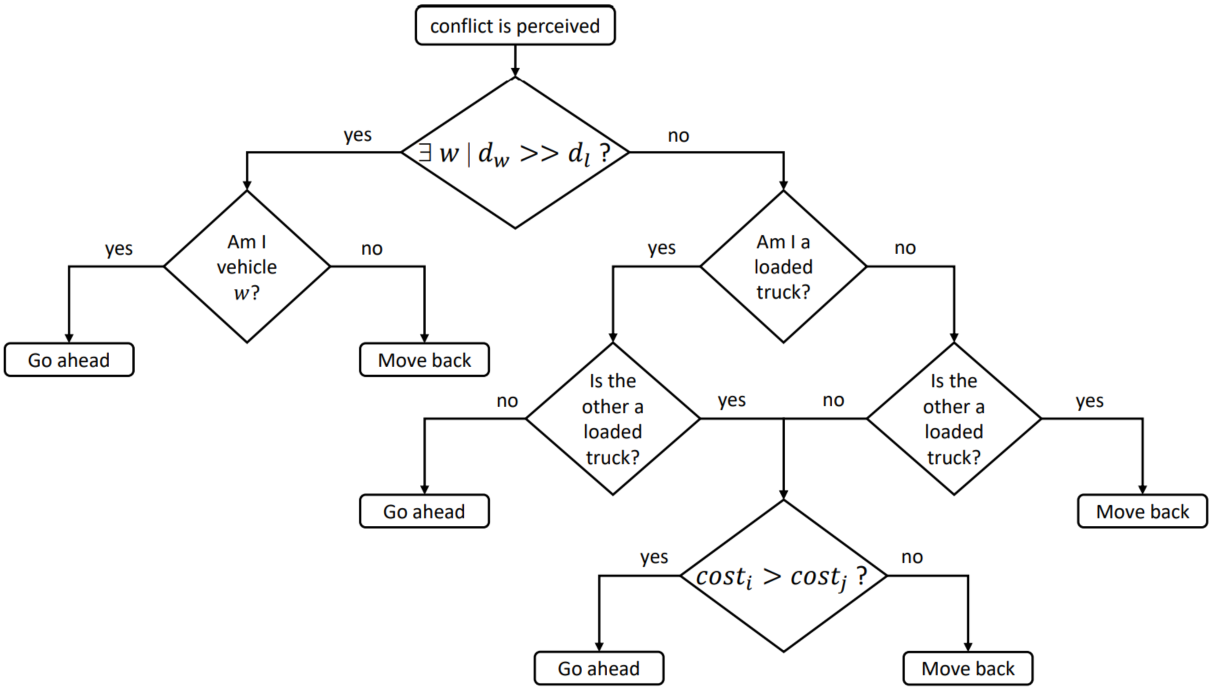

Figure 5.10 Behaviour of agents in the digital replica of the mine

Where:

- i is the agent who is making the decision, and j its counterpart in the conflict;
- w and l are the winning and losing vehicles of a trivial conflict in which there is a clear winner;
- d_v is the distance of vehicle v from the respective passing bay;
- $cost_i$ expresses how much it would cost for vehicle i to go into the passing bay and let the other one through. The author, in fact, decided to model the human decisions discussed above (lower part of the right-hand branch of the diagram in Figure 5.9) with a cost function. As mentioned earlier, in fact, real-world workers in certain situations choose the "winner" of the conflict by qualitatively considering certain criteria (distance from the bay; priority of the task). In the following approach, these approximate valuations made by the actors in the real mine are mapped into costs in the simulation model: each agent involved in the conflict will compute its own cost, and the one with the lowest value will step back to resolve the conflict. Since, in order to be true to reality, both the distances to the bays and the priorities of the current missions of the vehicles must be taken into account, the following cost function was designed:

  $$cost_i = \alpha\, cost_{i,distance} + (1-\alpha)\, cost_{i,priority}$$

  where:

  - ranging from 0 to 100, $cost_{i,distance}$ and $cost_{i,priority}$ express how much it would cost for vehicle i to move back in terms of distance and priority respectively. In the real mine these values are considered qualitatively: the author therefore considered it appropriate to devise a scoring system ranging from 0 to 100, to which each (digital) agent can refer in order to quantitatively assess the costs (which are qualitatively perceived in the reality). $cost_{i,priority}$, in particular, is equal to 0, or 20, or 40, or 60, or 80, or 100, for tasks with priority 0, 1, 2, 3, 4, 5 respectively (in the real mine, in fact, each activity is assigned a priority ranging from 0 to 5). $cost_{i,distance}$, on the other hand, is defined by the following formula:

    $$cost_{i,distance} = min(100\frac{d_i}{d_j}; 100)$$

  - Ranging from 0 to 1, α is the importance attributed to the distance criterion over the priority criterion. It is straightforward to realise that α is a function of d_i:

if the distance to the passing bay is a matter of a few metres, priority will have greater weight in the choice; if, on the other hand, the bay is hundreds of metres away, the distance criterion will dominate the priority one. The author, therefore, opted to model the relative importance of one criterion over the other with the following linear function (m denotes the slope):

$$\alpha = md_i$$

5.4.2 Simulation model calibration, validation and stability

The agent-based model outlined in the previous paragraph was implemented in AnyLogic8 university edition. A combination of agent-based and discrete-event modelling was employed to create the digital replica of the real mine, with the agent-based paradigm being predominant (please refer to Section 4.4).

With the digital replica at hand, the purpose is to exploit simulation as a decision support tool, by conducting experiments in a risk-free world. Before this can be done, it is necessary to make the simulation model operational, by conducting the validation and calibration phases: validation consists in checking that the digital replica actually behaves like the real-world system, while calibration consists in tuning certain model parameters for exactly the same purpose, i.e. to increase the accuracy of the model's representation of reality.

As it is easy to guess, for both phases it is necessary to identify some KPIs that are proxy indicators of the behaviour of the mine, be it the real-world one or the one simulated in AnyLogic. The model, in fact, will be validated and calibrated by comparing the KPIs values obtained in the simulated world with those obtained in the real mine: the greater the match between the values, the greater the fit of the model (i.e. the more the digital model behaves like the real mine). The behaviour of the mining logistics and production system, according to the author, can be seen as something that emerges from the dichotomy between its two complementary aspects: efficiency and effectiveness. For this reason, the chosen key performance indicators, required for the validation and calibration phases, relate on the one hand to the efficiency of the system, and on the other to its effectiveness. They are presented in Table 5.3.

While the efficiency KPIs will be qualitatively validated due to lack of data, mine effectiveness will be studied from a quantitative point of view. In any other simulation study, this would have meant comparing the KPIs value obtained from the simulation with that obtained in reality during the days that were simulated. Due to the high uncertainty inherent in underground mining, however, the author considered it inappropriate to follow this procedure. The

Table 5.3 KPIs to validate and calibrate the simulation model

Effectiveness KPIs	Efficiency KPIs
$T_i \quad \forall i = 1,2,3,4$ The mine extracts four types of minerals: waste, low-grade ore, high-grade ore, super-high-grade ore. Indicating the products as $i = 1,2,3,4$, T_i denotes the tonnes of product i extracted and transported to the surface	$Travel \quad Efficiency \quad (TE)$ This metric indicates how many kilometres the fleet travels to obtain one tonne of mined material, i.e. how much input must be provided to yield one unit of output. The reader interested in the formula and the meaning behind this KPI can refer to Chapter 7.2.1.
	$Resource \quad Utilisation \quad (U_f)$ The reader interested in the formula and the meaning behind this KPI can refer to Chapter 7.2.1.

next section will explain the why of this consideration and will present the countermeasures taken. The following paragraphs, then, will discuss the calibration, validation and stability of the conceived simulation model.

Benchmarking strategy

For research purposes, the simulation model of the mine was run to simulate one working day (consisting of two 12-hour shifts). The most straightforward and simple path to follow in order to calibrate and validate the model would have been to compare the value of the indicators found through the simulation with the value of the indicators obtained in the real mine on the day in question. As already hinted, the author considered this procedure not suitable for the world of underground mining, due to the high uncertainty at stake. Let us consider, for example, the indicator T_2 (the same applies to the other indicators) and the case shown in Figure 5.11.

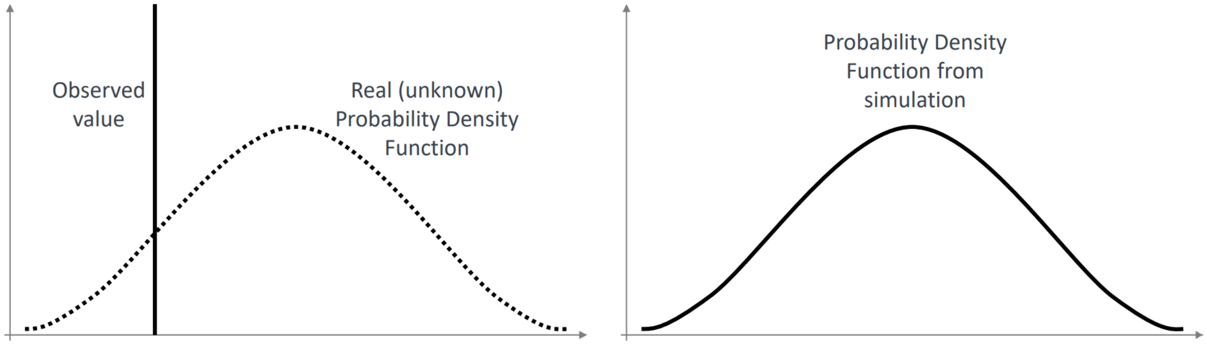

Figure 5.11 Real world (left) vs Simulated world (right)

Given the specific starting conditions (e.g. underground position of vehicles) and the production plan set for the day under consideration, T_2 (tonnes of low-grade ore extracted and transported to the surface after one working day) is a continuous aleatory variable with a given probability density function (on the left of Figure 5.11). The actual KPI value obtained by the mine after the working day in question is an observation of this aleatory variable. If it were possible to go back in time and work that day again, always with the same production plan and the same starting conditions, the found KPI value could be very different from the one observed the first time, due to the high uncertainty in the underground mine, which reflects into the high variance of the aleatory variable. The only data available, however, is the single observation of the random variable (continuous line in the left-hand side of Figure 5.11). The probability density function (dotted line in the left-hand side of Figure 5.11) is instead unknown: it is not possible to go back in time and work that day again as said before, and at the same time it is not viable to rely on past observations of the KPI, since the latter is a function of the production plan, which varies from day to day (the probability density function would therefore be tainted by the effect of the variation of production plans).

Let us now assume that the T_2 probability density function obtained through simulation (on the right of Figure 5.11) is exactly identical to the real (unknown) T_2 probability density function (ideal situation). This would mean that the simulation model is 100% valid, as it is able to reproduce exactly the same behaviour as the real system. If in the validation phase the aforementioned path was followed, the probability density function of the aleatory variable of the simulated world would be compared to the performance of the real system, which is nothing more than a simple observation of the aleatory variable of the real world (continuous line in the left-hand side of Figure 5.11). But, as already mentioned, due to the high uncertainty, which leads to the high variance of the aleatory variable, the available observation could be significantly far from the average. This would lead to the conclusion that the simulation model, which is actually 100% valid, is not valid. This type of error is known as *type I error*, and the related probability is called *model builder's risk* [47]. Concurrently, there is also evidence in the literature that the risk of accepting an invalid model, i.e. *type II error*, is dependent on the sample size [46]: the probability of this error, i.e. *model user's risk*, is therefore expected to be high as well.

The problem, therefore, stems from the fact that while in the virtual world it is possible to perform n replications of the 24 hours of work, and thus calculate the exact probability density function of the KPI in question, in the real world we only possess one observation of the aleatory variable, which has a high probability of not being well representative of the nominal behaviour of the system.

To solve the problem, the idea has been to define the behaviour of the real system not through the single observation from reality, but through its expected behaviour forecasted on the basis of its historical comportment. The use of m past observations, in fact, lends the comparison the statistical significance it lacked. In order to exploit past observations, however, the noise introduced by production plans must be taken into account, as mentioned above. For this reason, the author decided to study the percentage variation of the actual behaviour of the real system compared to the planned behaviour:

$$d_{ik} = \frac{a_{ik} - p_{ik}}{p_{ik}}$$

- i = i-th day considered
- k = type of product considered (waste, low-grade ore, high-grade ore, super-high-grade ore)
- a_{ik} = actual value of production of k in the i-th day
- p_{ik} = planned value of production of k in the i-th day (defined by the production plan)
- d_{ik} = percentage variation of production of k in the i-th day

For each product type, a dataset on the 84 days preceding the simulated one was analysed. For each k, therefore, d_{ik} was calculated, with i ranging from 1 (first day considered) to 84 (day preceding the simulated one). In this way, for each product type, it was possible to study the aleatory variable D_k, which represents the percentage variation of the actual production compared to the planned one in a generic day. The 90% confidence interval for the mean of D_k was then calculated for each product: $[l_{D_k}, u_{D_k}]$ is the 90% confidence interval for the mean of D_k. Let us now focus on the simulated day, i.e. the 85th day. The production plan is known, so p_{85k} is known for every k. Assuming that the mine behaves as it has done in the past, it is possible to state with a 90% confidence level that the expected value of D_k is between l_{D_k} and u_{D_k}, known for each product k. From this it can be deduced that the expected value of the actual production in the 85th day, again with a confidence level of 90%, is between $p_{85k} + p_{85k}l_{D_k}$ and $p_{85k} + p_{85k}u_{D_k}$ for each product k. For each product we have therefore found a range of values within which the actual production on the day considered should lie if the mine behaved as it did in the past 84 days, with 90% confidence. Figure 5.14 shows these intervals (labelled with the term *expected*). The production values obtained from simulation, in conclusion, will be compared with these ranges in order to calibrate and validate the model.

Model Calibration

Once the simulation model was set up with AnyLogic8 university edition, descriptive statistics was used to map real-world observations into model parameter values. It was not possible, in this regard, to directly measure two parameters: m and *visibility distance*. m is the slope of the line $\alpha = md_i$, where α represents the weight attributed by the driver to the criterion of distance over that of priority when deciding how to resolve a conflict, and d_i is the distance of the vehicle from the nearest passing bay. For a better understanding, the reader is invited to refer to the formulas already presented in Section 5.4.1. m, and consequently α, are parameters of the simulation model with which the author has sought to model the decisions made by humans in the real system, and which cannot therefore be directly measured. The *visibility distance* parameter, on the other hand, is linked to the agents' perception of their environment. In the case of head-on conflicts, in which two vehicles drive through the same underground tunnel in opposite directions, each agent perceives the danger only when a certain distance from the other, i.e. the *visibility distance*, is reached. This value in the real mine varies from conflict to conflict, as it is influenced by factors specific to the individual situation, such as the nature (curved or straight) and lighting conditions of the stretch of road in question. Unlike all the other parameters of the simulation model, therefore, it was not possible to define a value for m and *visibility distance* a priori. In order to make the simulation model operational and ready for the planned experiments, therefore, it was necessary to determine the value of these two parameters. The best way to do this is to use calibration, a process that exploits data on the behaviour of the real system and fine-tunes the values of parameters to increase the goodness of fit of the simulation model to reality. AnyLogic calibration experiments were used in this thesis to calibrate the two parameters. They were obtained by minimising the deviation between the values of the effectiveness KPIs (proxies for behaviour) found with the simulation and the centre of the desired ranges of KPIs values (shown in Figure 5.14). The parameters m and *visibility distance* were hence tuned in order to increase the match between the behaviour of the simulation model (how the digital replica behaves) and the desired behaviour (how we want the digital replica to behave).

Model stability

In an underground mine, operations that are supposed to take minutes may actually take hours, the same journey between two underground points may have very different durations if done several times, and so on and so forth. The conceived simulation model, as a consequence, is characterized by a high degree of stochasticity. For proper calculation of statistical data and in order to produce statistically significant results, therefore, it was necessary to replicate several times the simulation experiment in AnyLogic. This section deals with the choice of the number of replications and the stability of the model.

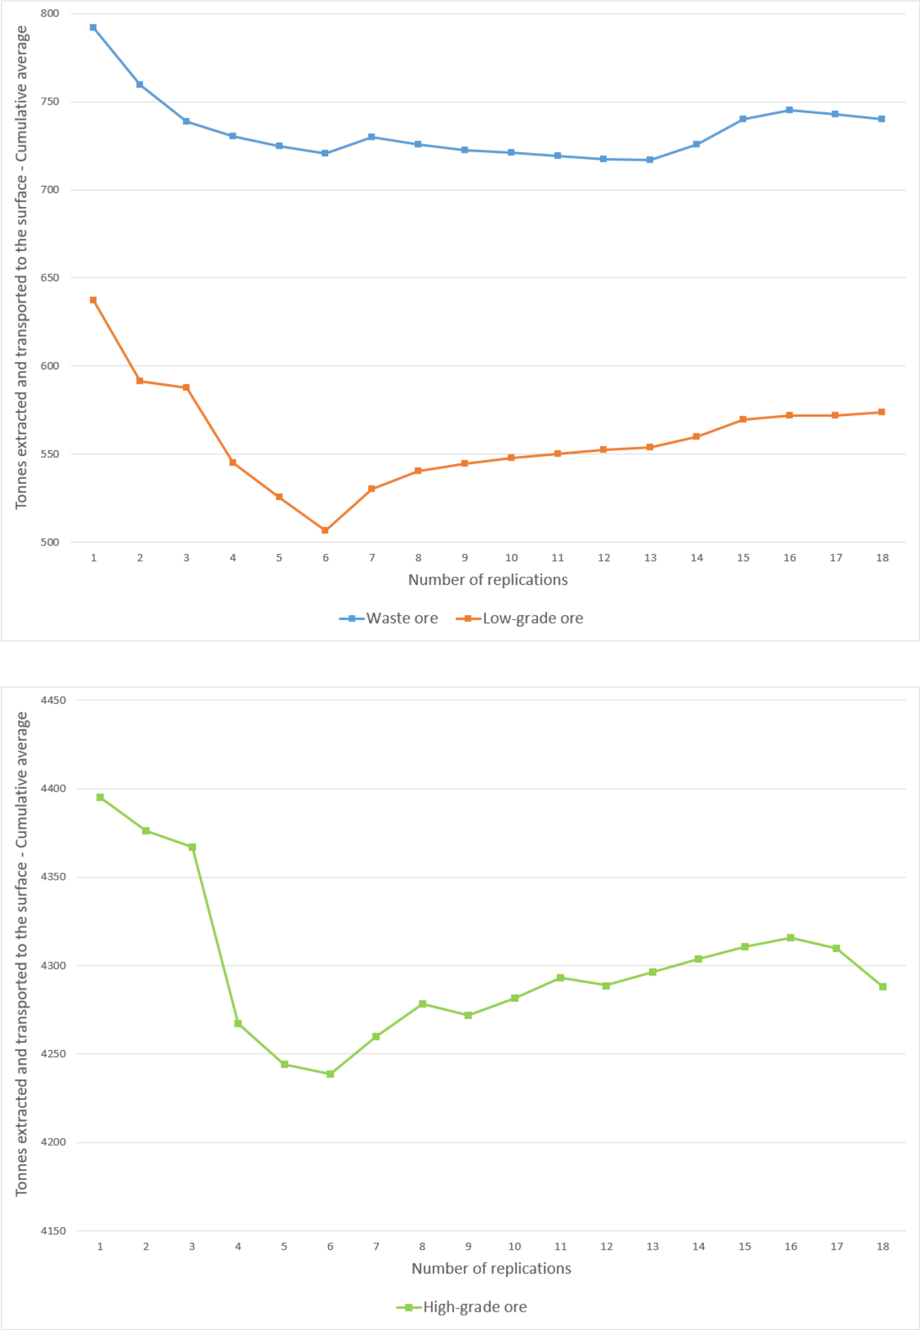

Figure 5.12 Simulation model stability

The graphs in Figure[2] 5.12 present, for each number of replications of the simulation experi-

---

[2] On the simulated day, the mine does not plan any quantity of super-high grade ore to be extracted: for this reason, this type of ore will not appear in the results that will be presented in the various sections of the thesis.

ment, the cumulative average for the effectiveness KPIs. By way of example, the cumulative average of a KPI corresponding to 6 replications is equivalent to the arithmetic mean of the 6 KPI values observed with the 6 replications. As shown in Figure 5.12, it appears that the cumulative averages of the main KPIs tend to stabilize after 10 simulations. After this value the simulation model becomes stable, i.e. the addition of a further replication does not lead to significant changes in the average value of the main KPIs. The average values obtained by means of 10 replications will therefore be employed for the analyses conducted in the rest of this thesis.

Model validation

Verification and validation processes play a crucial role when developing a simulation model. Verification serves to ensure that the conceptual model has been transformed into a computer model with sufficient accuracy [48], while validation is usually defined to mean "substantiation that a computerized model within its domain of applicability possesses a satisfactory range of accuracy consistent with the intended application of the model" [49]. In simpler words, validation is about building the right simulation model while verification is about building the simulation model in the right way. Over the years, several approaches have been proposed for validation and verification in simulation studies. Most of them elaborate on and boils down to the concepts proposed for the first time by [4]. For this reason, the paradigm proposed by [4], presented in Figure 5.13, was followed in this thesis.

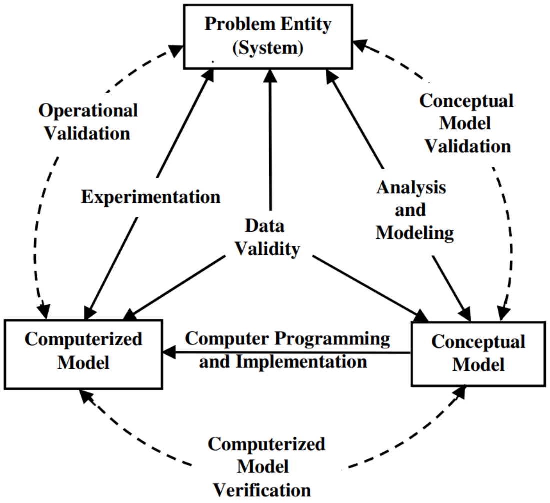

Figure 5.13 A simplified version of the modelling process by [4]

At this point in the discussion, the three main elements of the paradigm have already been discussed: the *system* was presented in Sections 5.2 and 5.3, which dealt with the industrial partner logistics and production system and the related management approach; the *conceptual model* was illustrated in detail in Section 5.4.1, which presented the agent-based model of the underground mine; the conceptual model was finally implemented in AnyLogic8 university edition by means of Java programming language, giving rise to the *computerized model*. If the reader is interested in knowing more about the other individual aspects of the paradigm, he/she can refer to [4]. The goal of the rest of this section, in fact, is to delve into *operational validation*, since it is of paramount importance in the process: it determines if the model's output behaviour has sufficient accuracy for the model's intended purposes, i.e. it informs us whether the created simulation model is "correct". Unfortunately, there is no standard procedure to follow to determine the "correctness" of a simulation model, and there is no reference guide to determine which techniques to use each time. Each simulation project presents a new and unique challenge to the model development team. [46], however, proposed several viable paths for operational validation, and is one of the most widely used references when it comes to validation and verification of a simulation model. According to [46], when dealing with observable systems (i.e. when it is possible to collect data on the system behaviour), as in the case of this thesis, two main approaches can be followed. The first approach consists of exploring the model behaviour: this means validating the directions and magnitudes of the model's output behaviour when, for example, conducting a sensitivity analysis on the parameters. In this thesis, for instance, it was investigated how varying the speed of vehicles affected the effectiveness KPIs in terms of direction (increase or decrease in the total amount of material produced and transported to the surface) and magnitude (extent of the increase/decrease). The second approach is the so-called comparison: it means comparing the simulation model output behaviour to the real system output behaviour using graphical displays and/or statistical tests and procedures. In this thesis confidence intervals have been employed to compare the behaviour of the digital replica with that of the real mine (Section 5.4.2). Figure 5.14 below also graphically shows the results of the benchmarking, highlighting the deviation between the real and simulated values of the KPIs used to define the behaviour of the mine.

The figure, in particular, compares for each type of extracted material:

- the value ranges related to the behaviour of the real-world mine (*expected*); these ranges reflect the expected behaviour (90% confidence) of the real mine forecasted on the basis of past performance (see section "Benchmarking strategy");
- the behaviour of the digital replica (*simulated*); figures refer to average values from 10 simulations.

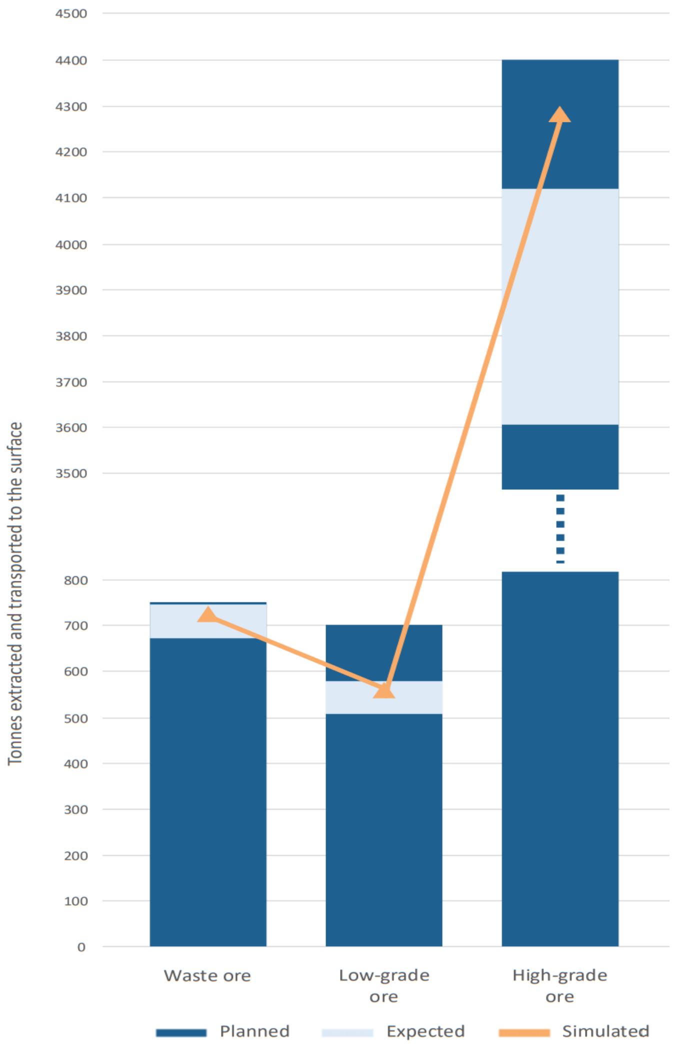

Figure 5.14 Simulation model validation

The comparison suggests that the created simulation model has a good degree of accuracy in representing reality. The values obtained with the simulation are in fact within the desired ranges. The only exception is the high-grade ore, in relation to which the digital replica seems to over-perform the real mine. It must be considered, however, that the mine created in AnyLogic belongs to an ideal world in which some disturbances (such as workers' restroom breaks and unexpected events) are not taken into account: it is therefore acceptable to obtain production levels slightly higher than the real ones. The model created thus seems to have an acceptable "goodness of fit".

[46] also discussed practical approaches to verification and validation of simulation models. Table 5.4 lists the techniques presented by [46] that have been employed in this thesis, providing a brief description and an example for each of them.

Table 5.4 Validation techniques

Validation Technique	Description	Application in this thesis
Animation	The model can be visually validated thanks to the graphic display of its operational behaviour	A 3D animation of the mine was embedded in the simulation model built with AnyLogic
Operational Graphics	"Values of various performance measures are shown graphically as the model runs through time; i.e., the dynamical behaviors of performance indicators are visually displayed as the simulation model runs through time to ensure they behave correctly" [46]	When launching the simulation, in parallel with the 3D animation of the mine, the user has access to interactive and immersive dashboards that provide real time information on the main KPIs (the interested reader can find the dashboards in Appendix B). An analysis of the KPIs dynamic behaviour was conducted to validate the model
Traces	"The behaviors of different types of specific entities in the model are traced (followed) through the model to determine if the model's logic is correct and if the necessary accuracy is obtained" [46]	During the simulation, AnyLogic offers the possibility to analyse the real-time state and behaviour of each agent. In the digital replica of the mine, the author also installed a camera on board each virtual vehicle, which allows the user to immerse into the simulation by climbing aboard the desired car (Appendix B shows an example of climbing into a vehicle during simulation). Both tools proved to be very useful in validating the individual behaviour of the model agents
Degenerate Tests	"The degeneracy of the model's behavior is tested by appropriate selection of values of the input and internal parameters" [46]	The duration of operations in underground working areas was made to skyrocket: queues of resources were observed to form in the working areas, with the average number of vehicles in the queue continuing to increase over time
Input-output validity	"Compare the model input-output transformations to corresponding input-output transformations for the real system" [46]	Using the same production plan and dispatching lists as in the real mine, the simulation model was able to produce an output quite similar to the actual figure (Figure 5.14)
Parameter Variability - Sensitivity Analysis	"This technique consists of changing the values of the input and internal parameters of a model to determine the effect upon the model's behavior or output." [46]	A sensitivity analysis was conducted around the speeds of vehicles

To summarise what has been said so far, the devised simulation model was validated by: exploring the behaviour of the model; comparing its behaviour with that of the real mine, using both graphical displays and statistical procedures; exploiting the techniques shown in Table 5.4. Conducting iterative validation tests of various nature and adjusting the model from time to time according to the findings allowed the author to attain a model with a good degree of accuracy with respect to the intended purposes of the simulation study. Stakeholders in the simulation project have thus nurtured a high level of confidence in the conceived simulation model and its results. This achievement was a milestone for the research project. In the remainder of this work, in fact, the author will aim to improve the mining system and the related management practices by exploiting multi-agent systems. The simulation model of the *system to be*, i.e. the mine incorporating the envisaged improvements, will be built from the simulation model of the *system as is* just presented. Having a valid starting model, therefore, was of utmost importance for the success of the rest of the project.

CHAPTER 6 PROPOSED AGENT-BASED FLEET MANAGEMENT SYSTEM

6.1 Introduction

This chapter is structured as follows. First, the backbone of the proposed solution is presented in Section 6.2. Next, the specific fleet management strategy is outlined. Finally, Section 6.4 integrates the two aspects of the two previous sections, presenting the overall solution proposed by this thesis.

6.2 Backbone: agent-based systems and Industry 4.0

Oriented by the vision of the smart factory envisaged by Industry 4.0, this thesis proposes an agent-based fleet management system, in which software agents make dispatching, routing and scheduling decisions. Each vehicle in the fleet is controlled by a software agent which is physically deployed on the car, just as each underground site is outfitted and controlled with the related software agent. Company's resources and underground working areas are thus defined as autonomous agents working together towards a common goal. The resulting system is computerized and composed of multiple interacting autonomous entities. In the proposed solution, then, smart sensors installed in the underground mine track the real-time position of each vehicle, and serve as the interface between the digital (agent-based) and physical worlds. Each vehicle is also equipped with an on-board navigation system, through which the associated software agent provides information to the driver about the next destination to visit (dispatching decision), the route to follow in case of alternatives (routing decision), and traffic coordination with other vehicles (scheduling decision). In a nutshell, while in the as-is mining system dispatching decisions are made by mine engineers at the beginning of the work shift and routing and scheduling decisions are delegated to vehicle drivers, in the to-be mining system software agents take over the control of the three decision-making processes. They take dispatching, routing and scheduling decisions by exploiting the real-time information provided by the sensors, and report the results of the decision-making to vehicle drivers via the on-board navigation systems.

With the described cyber-physical system in place, specific fleet management approaches can be implemented, i.e. software agents can be fine-tuned to execute tailor-made management strategies. Section 6.3, in this regard, will introduce the reader to the management approach proposed by this thesis, illustrating how software agents controlling vehicles and underground sites behave and coordinate to make dispatching, routing and scheduling decisions.

The design process from which the presented ideas came up was oriented by a design vision centred on the smart factory envisioned by Industry 4.0. This paradigm has therefore nurtured and provided a sense of direction for this thesis. This research project, as a result, can be seen as a first step on the path towards the mine 4.0. First and foremost, in fact, at the heart of the Industry 4.0 philosophy are cyber-physical systems, just like the mining system proposed in this thesis, which springs from the dichotomy between the digital world in which the software agents live and the physical world inhabited by the real-world vehicles, workers and sites. Boston Consulting Group [50], moreover, defined nine main enabling technologies of Industry 4.0, among which is the Industrial Internet of Things, which forms the backbone of the designed system. Software agents, smart sensors and on-board navigation systems, in fact, all embrace this paradigm, which envisages the use of multiple inter-connected devices and sensors exchanging manifold information about production and logistics processes and company's assets, allowing for data collection, exchange, and analysis. One of the main principles underpinning Industry 4.0, to conclude, is to digitise information and ensure its transparency, in real time, in order to support the decision-making process and promote its decentralisation. In the proposed solution, in this regard, software agents of the agent-based system use real-time information provided by sensors to make dispatching, routing and scheduling decisions in a decentralised, distributed and autonomous manner, interacting and collaborating with each other in order to achieve the common goal. The main vision behind the Industry 4.0 paradigm, always referring to fleet management, is therefore pursued with the multi-agent system presented in this thesis.

Last but not least, a remark must be made about the level of automation of the proposed mining logistics and production system. At this stage of industrial evolution, in fact, there are still many tasks in production and logistics processes that are not well fitted for fully automated machines and for which human workers are still a key resource. In compliance with this status quo of Industry 4.0, therefore, the designed solution features a medium level of automation: although management decisions are made by software agents, it is then up to the workers to execute the output of the decision-making process and to drive the vehicles, thus still playing an essential role in the process. It must be emphasised, however, that the conceived agent-based system is still valid in the case of mining systems featuring a higher level of automation. The solution would work for mines in which most of the underground work is automated, no one works near the front anymore, and vehicle control takes place from collaborative control rooms above ground. The solution would also work for fully-automated mines featuring self-driving mining vehicles and in which drivers are completely outside the logistics process. The mining system presented in the rest of the thesis, therefore, is meant for a mine at the early stages of the path towards the fully automated mine, but is perfectly adaptable to the increasing degree of automation that will be encountered along the way towards the smart mine envisaged by Industry 4.0.

So far the backbone of the proposed solution, grounded on Industry 4.0 and agent-based technology, has been outlined. The next sections will shape the comprehensive fleet management system: while so far we have explained “who” is in charge of fleet management decisions, i.e. the software agents of the agent-based system, in the next sections we will focus on “how” agents make these decisions (Section 6.3), and on the specific design of the multi-agent system (Section 6.4).

6.3 Fleet management strategy

This thesis deals with fleet management in underground mines. Such a management problem consists of making, for each vehicle in the fleet, three operational decisions:

1. Dispatching decision
   Assign the next destination to a resource that has just finished its current task.
2. Routing decision
   Define the route to be followed to reach the chosen destination point.
3. Scheduling decision
   Define the vehicle movement so as not to have traffic conflicts with other vehicles moving in the mine. This decision is necessary because of the nature of the transportation network, which consists of one-lane bi-directional road segments.

The proposed solution is an agent-based system in which the three management decisions are delegated to software agents embedded in each mining vehicle and in each working area. Smart sensors allow the system to be operational by providing information to software agents, while on-board navigation systems are the interface through which decisions taken by the agent-based system are communicated to vehicle drivers. The resulting mining system is a cyber-physical system, which enables the mine to take the first steps towards the smart mine envisaged by Industry 4.0. All that remains to be done is to understand how the software agents of the multi-agent system make dispatching, routing and scheduling decisions. In the remainder of this section, therefore, the reader will be presented with the designed fleet management strategy, disentangling it into the dispatching, routing and scheduling decision-making processes. We can anticipate that the proposed dispatching strategy is a radical innovation on the one currently adopted by the mine under investigation, while the routing and scheduling strategies can be seen as incremental improvements.

Dispatching strategy

Considering a resource that has just finished its current task, making the dispatching decision means choosing what its next assignment will be, thus answering the question “where should it go now?”.

To solve this task assignment problem, the proposed solution is a Dynamic Contract Net Protocol (DynCNET), tailor-made for underground mines, which allows the implementation of dynamic multi-agent scheduling strategies.

DynCNET was pioneered by [51] to solve the same and identical problem addressed in this thesis in the industrial world of AGVs: it was proposed as a tool for flexible transport assignment in AGV transportation systems. DynCNET, more specifically, is the dynamic and flexible version of the Contract Net Protocol [52], which has been widely employed in the literature to allocate tasks among autonomous agents in agent-based systems. Since this research project proposes exactly an agent-based system as a solution to the fleet management problem in underground mines, the author considered CNET and its dynamic version to be appropriate and profitable tools to solve the dispatching problem. The results obtained by [51] applying DynCNET in an agent-based system designed for AGV transport systems also encouraged this work, since the fleet management problems in underground mines and in the world of AGVs present evident similarities: the transportation system, its environment (composed of one-lane bidirectional road segments), and the decisions to be taken to manage the fleet are quite similar in the two contexts (the reader can refer to Section 3.2 if interested in exploring these parallelisms).

To understand the functioning logic of the proposed solution to dispatching, first a CNET protocol and then its dynamic version, both tailored to the world of underground mining, will be presented.

CNET [52] is a coordination mechanism for task assignment, which prefer coordination through interaction protocols, and which stresses the utility of negotiation as an interaction mechanism. An agent (the Initiator) assumes the role of manager who wishes to have a certain task performed by another agent. It therefore contacts with a call for proposal a set of agents (the Participants), which assume the role of contractors. They can decide whether to submit a bid or not. Among the received proposals, the initiator chooses the best one by optimising a function that characterises the task. Exploiting market mechanisms, therefore, the task is allocated to a contractor. This is the rationale of the CNET protocol. With the general idea in mind, we can now discuss the Contract Net Protocol designed by the author to solve the dispatching problem in underground mines, illustrated in Figure 6.1.

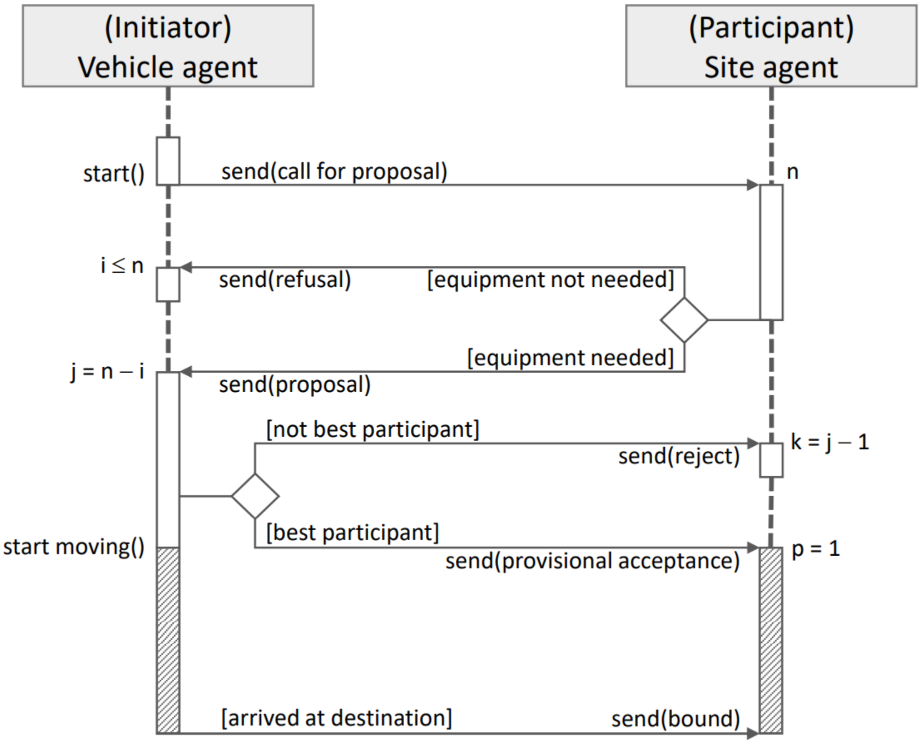

Figure 6.1 Designed CNET protocol

Before explaining how the protocol works, it is worth highlighting a design choice made by the author. CNET, in particular, is an $m \times n$ protocol: an initiator offering a task can interact with m participants, i.e. candidate agents that can perform the task, while each participant can interact with n initiators offering tasks [51]. In the conceived agent-based system, the most obvious choice would have been to have the underground sites, which "generate" the tasks, play the role of initiators, and the vehicles, which are able to perform the tasks, play the role of participants. In an underground mine, however, mining resources form the bottleneck of the system and the number of tasks to be performed is much greater than the number of resources capable of executing them. This means that, if the aforementioned design choice were adopted, n would be much greater than m. This is why the author has chosen to cast the sites in the role of participants and the mining vehicles in the role of initiators. This decision, however, has been a design choice: it would be interesting to study in future works the performance of the same DynCNET protocol, in which working areas and mining resources swap the roles of participants and initiators.

Having understood the roles of the various types of agents in the protocol, we can now move on to the functioning of the CNET illustrated in Figure 6.1.

Whenever a mining equipment completes its current task, it offers its availability to the mine with a call for proposals addressed to all underground sites. It therefore initiates contract negotiation by advertising the existence of its availability. At this point all sites interested in the services offered by that resource will place an offer. This means that the bid will be submitted by all sites that are in a state of the development cycle for which that resource is required (e.g. if the equipment in question is a charger, all sites that have already been drilled and now require to be loaded with explosives will bid). Among the received proposals, the initiator chooses the best one. Two factors, in particular, need to be taken into account to assess the goodness of an offer: the priority index of the applicant site and its distance from the current position of the vehicle. In this way one tries to be efficient (minimising the distance to be travelled), while still respecting the constraints imposed by the production plans (accounted for through priority indexes). At this point in the negotiation, a destination has been assigned to the resource, i.e. the dispatching decision has been made. In the traditional CNET protocol, the coordination between the agents of the agent-based system would end in this way, the vehicle would reach the assigned destination and perform the corresponding task. In order to cope with the high uncertainty and dynamism of underground mines, however, this protocol is going to be made dynamic and flexible, giving birth to the DynCNET: the assignment resulting from the classic CNET negotiation will no longer be definitive, but will simply be a provisional agreement between resources and sites, which can be modified at any time before the vehicle arrives at destination, in order to benefit from emerging opportunities and mitigate emerging threats in the mine. Shaded areas in Figure 6.1, in this regard, represent time windows in which agents can switch the provisional agreement. In the time frame from the creation of the temporary agreement to the arrival of the vehicle at destination, therefore, both parties involved in the deal have the possibility to withdraw from the contract, in order to take advantage of (mitigate) dynamically arising opportunities (threats) that were not present when the provisional contract was stipulated. In such a case, the party withdrawing from the contract does so in order to engage in a more convenient temporary agreement with a third party, while its counterpart will have to start again from the early stages of the protocol and find a willing fourth party. As soon as a resource engaged in an open provisional agreement reaches its destination, in any case, it notifies the counterpart (underground site) with a bound message: the temporary agreement becomes an actual contract binding the vehicle agent and the site agent, and the task can finally start. This is the general logic behind the DynCNET protocol, whose details will be presented more specifically in the following.

The dynamic component that is incorporated into the CNET protocol in Figure 6.1, hence giving rise to DynCNET, can be presented with the help of the statechart diagrams in Figure 6.2.

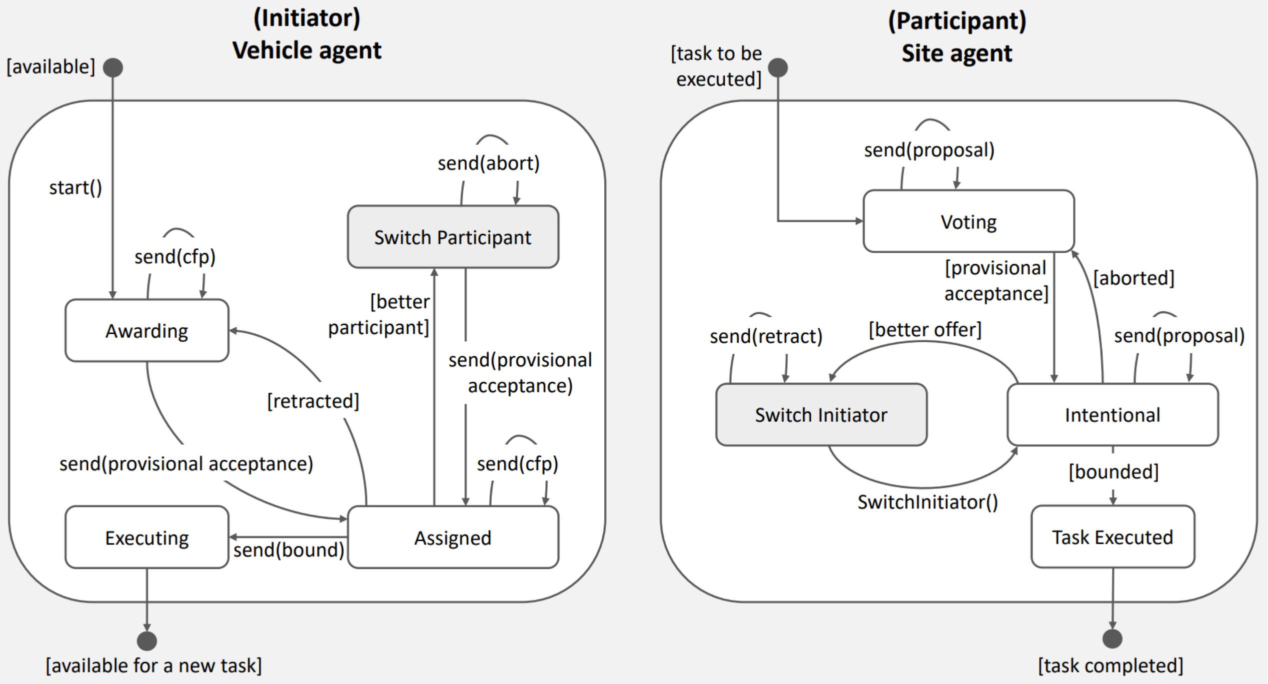

Figure 6.2 Statecharts of the agents

The two statecharts reveal the behaviour of the agents participating in the negotiation of the DynCNET protocol. They highlight, in particular, how agents behave in order to modify, if appropriate, the provisional assignment stipulated via the "classic" CNET protocol, in order to take advantage of opportunities and mitigate threats that dynamically emerge in the underground mine. Agents, in particular, switch the provisional agreement in the shaded states of the statecharts. Figure 6.2, to conclude, will be explained using two points of view: the one of the initiator who switches participant and the one of the participant who switches initiator.

Switching Participant

When a resource finishes a task and wonders where to go next, it advertises its availability to the whole mine through a call for proposals. It evaluates the bids received from site agents interested in its services and chooses the best one considering distances and priority indexes. Sending a provisional-acceptance message to the "winner of the auction", it commits to a temporary agreement, entering the *assigned* state and starting to move towards the counterpart. The peculiarity of DynCNET lies in the fact that such an assignment is provisional: even if the dispatching decision has already been made, i.e. a new task has been established for an available resource, this decision is reconsidered dynamically over time. Indeed, despite having already been assigned, the initiator continues to send out calls for proposals to the entire mine, behaving de facto as an agent not yet assigned. While heading towards the counterpart of the temporary contract, therefore, if a better site (always considering both priority indexes and distances) becomes available, the initiator switches the participant. A withdrawal message will be sent to the counterpart of the open contract, that will thus have to start participating in new auctions again (back to *voting* state) to find a new resource able to treat it. The initiator, on the other hand, stipulates a new, more advantageous temporary contract with a new site, to which a provisional acceptance message is sent. This behaviour is iterated in time, until the vehicle reaches the counterpart of the current open temporary contract: at this point the temporary agreement becomes an actual commitment, the initiator sends the notification to the site and starts to execute the task.

An example of a switch of participant is shown in Figure 6.3, where a charger on its way to the site A, with which it has a temporary contract, reconsiders the dispatching decision: site B has just finished being drilled and represents a better match for the charger (because it has a higher priority than A and is closer to the charger), which therefore backs out of its commitment to A for a more advantageous deal.

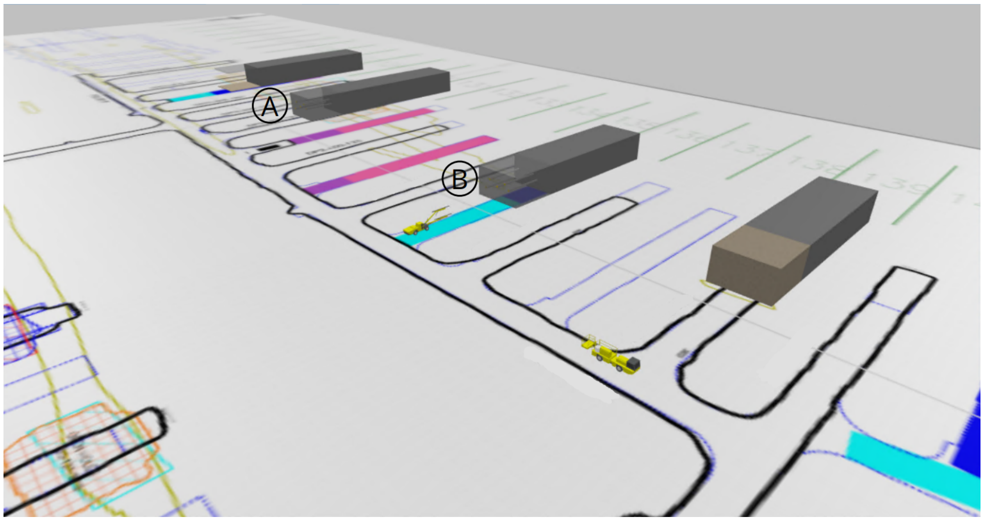

Figure 6.3 Example of switch of participant

Switching Initiator

An underground site is iteratively subjected to development cycles, i.e. ordered sequences of activities, each conducted by the corresponding resource (Figure 1.1). A site agent needing the execution of a given activity, therefore, participates in auctions set up by all mine resources capable of performing that activity. When an auction is won, the agent receives provisional acceptance from the auctioneer resource and enters the *intentional* state. Similarly to the behaviour of the initiator, the participant has the possibility to opt out of the deal: although there is already a resource willing to treat it and which, as per agreement, is heading towards it, the agent (in the *intentional* state) continues to participate in auctions of other resources by submitting proposals. If a better initiator (i.e. a vehicle capable of performing the task closer to the site than the vehicle with which a temporary contract has been concluded) is found, the agent switches initiator and stipulates a more favourable agreement. In this case, a withdrawal message is sent to the old initiator, who will thus return to the *awarding* state and offer again its availability to the mine with calls for proposals. The site agent will continuously monitor the presence of new opportunities in the mine, until the vehicle with which the last temporary contract was established arrives at destination, sending a bonding message and initiating the task.

An example of a switch of initiator is shown in Figure 6.4, where the site X, despite having already entered into a provisional agreement with driller A, signs a more advantageous contract with driller B, that has just finished its task at site Y. Driller B represents in fact an opportunity for X, since it is closer to X than driller A is. Site X, therefore, revokes the contract with A and creates a new one with B.

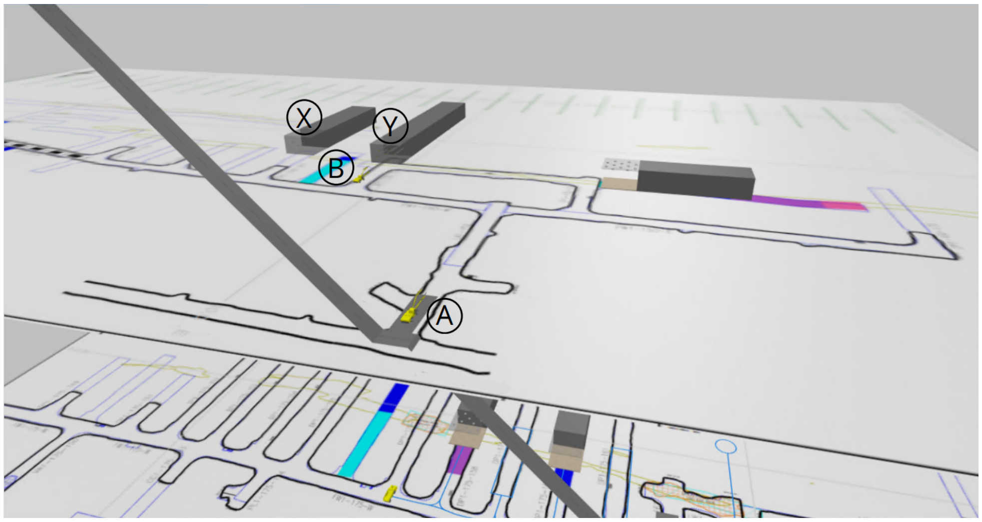

Figure 6.4 Example of switch of initiator

To summarise what has been said so far, the solution proposed by this thesis to the dispatching problem in underground mines is the DynCNET protocol, a dynamic multi-agent scheduling approach. By means of this coordination mechanism between the agents of the proposed agent-based system, the logistics and production system gains a high level of agility, which allows it to cope with the high level of uncertainty and dynamism inherent in the world of underground mining. The mining system, therefore, will be able to rapidly exploit information from its dynamic and unpredictable environment, and to adapt in order to efficiently and effectively achieve its objectives. The results achievable using DynCNET and multi-agent dynamic scheduling will be estimated by implementing them in the digital replica of the industrial partner created with AnyLogic. A quantitative analysis of the results obtained with the simulation study will be presented in Chapter 7.2.2.

Routing strategy

While the envisaged dispatching strategy based on DynCNET and dynamic multi-agent scheduling represents a radical innovation compared to the approach currently used by the industrial partner, the routing strategy presented below can be seen as an incremental improvement.

The routing decision consists of choosing which path to follow when there are several alternatives leading to a same point in the mine. While in the system "as is" this decision is delegated to the drivers of the vehicles, who, under normal operating conditions, choose the shortest of the alternative routes, in the system "to be" (an agent-based cyber-physical system) the decision will be entrusted to the software agents controlling the vehicles. It will therefore be possible to exploit the advantages offered by automation, and take "better" routing decisions by making software agents do calculations that humans would not be able to do. Mine workers, in particular, when choosing the shortest route have no way of taking into account the traffic conditions on that road, due to the limited visibility in the mine underground tunnels. The shortest route in terms of distance, therefore, may turn out to be the longest in terms of travel time. In the proposed system, vice-versa, software agents interacting with sensors are able to know at any time the number of vehicles on a road, and thus get an idea of traffic conditions. The insight, therefore, is to make the routing decision on the basis of both the distance factor, as done in the as-is system, and the traffic factor.

The resulting decision-making process, which unfolds in the behaviour of vehicle agents, can be described with the help of the following pseudo-code:

Algorithm 1: Routing decision

```
1  routes $\leftarrow$ routesAvailable;
2  routeChosen $\leftarrow$ null;
3  while $routeChosen = null$ do
4      shortest $\leftarrow$ SHORTESTROUTE(routes);
5      if $TRAFFIC(shortest) < \gamma$ then
6          routeChosen $\leftarrow$ shortest;
7      else
8          routes $\leftarrow$ routes $\setminus$ shortest;
9          if $routes = \varnothing$ then
10             routeChosen $\leftarrow$ SHORTESTROUTE(routesAvailable);
11         end if
12     end if
13 end while
14 return routeChosen;
```

where:

- *routesAvailable* is the collection of routes available to reach the destination point from the point of departure.
- *SHORTESTROUTE(routes)* returns the shortest path in the collection of routes called "routes"
- *TRAFFIC(route)* returns the traffic intensity in the given route; the traffic density is measured by dividing the number of vehicles present on the road by its length, and is therefore expressed in $\left[\frac{vehicles}{meter}\right]$
- γ is the threshold above which the level of traffic is deemed unacceptable

The pseudo-code outlines how a software agent controlling a vehicle, given a set of available roads: considers the shortest route as the best candidate, but before selecting it as the "winner" assesses whether its traffic density is deemed acceptable; if so, the road is selected and the routing decision is made; if not, the agent considers the second shortest route in the list and iterates the procedure. The code in lines (9) and (10) indicates that in case all roads in the list are too crowded, i.e. the traffic is deemed unacceptable for each of them, the shortest route is selected. Using another perspective, we can say that software agents, in order to decide which road to take among the alternatives, solve an optimisation problem in which the objective function to be minimised is the distance to be travelled, and the traffic density is a constraint to be satisfied.

Scheduling strategy

As with the routing strategy just presented, the solution devised for the scheduling part of the problem can be seen as an incremental improvement of the approach currently employed by the mine under investigation.

The scheduling strategy will be introduced by means of the example in Figure 6.5.

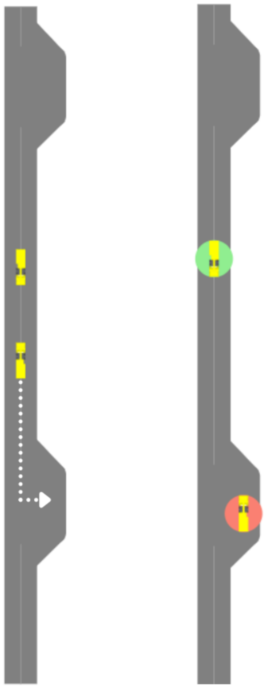

Figure 6.5 Scheduling strategy: “as is” (left) vs “to be” (right)

As one can observe, the underground transportation network consists of one-lane bidirectional tunnel segments, with the presence of occasional passing bays where vehicles can temporarily stop to let others pass. Please note that with the term “tunnel segment” the author is referring to the stretch of road between two successive passing bays. When moving in the underground mine, therefore, mining vehicles must coordinate with each other to cope with the nature of the transportation network.

In the as-is system (left-hand side of Figure 6.5), the decision of how to coordinate in traffic is delegated to the drivers of the vehicles. Two cars travelling in opposite directions perceive the existence of the head-on conflict when a visibility distance is reached, i.e. when both vehicles are able to see the counterpart (note, we are talking about underground tunnels: visibility is therefore limited). Upon sensing the conflict, one of the two backs up to the passing bay, lets the oncoming vehicle pass, and finally resumes its journey.

The idea is to eliminate the back-and-forth journey of the vehicle which goes back to the bay to resolve the conflict. Such a movement, in fact, is a non-value-added activity. Inspired by the Lean Manufacturing philosophy, which, among other things, focuses on reducing all non-value-added activities, and in which transportation is one of the *seven deadly wastes*, the author therefore designed an improved approach to scheduling.

In the proposed solution (right-hand side of Figure 6.5), software agents controlling vehicles, supported by the sensors installed in the mine, coordinate using the First In First Out (FIFO) logic: the first agent to reach the tunnel will be the first to be "processed", i.e. to occupy the tunnel. The software agent of the first vehicle to enter a tunnel segment (i.e. the stretch of road between two successive bays) will be informed by sensors that the segment is free: the on-board navigation system will show a green light to the driver, who will continue driving and enter the tunnel. The software agent of the second vehicle wishing to enter the tunnel and coming from the opposite direction will be informed by sensors that the segment is already occupied: the on-board navigation system will show a red light to the driver, who will stop in the passing bay, wait for the passage of the vehicle coming from the opposite direction, and finally resume its journey. This is the general idea behind the interaction mechanism between agents, which was then further detailed in order to consider more than two stakeholders, again based on the FIFO rule.

Delving into more detail, when a vehicle approaches the entrance of a tunnel segment, it must decide whether to enter the tunnel or wait in proximity of the passing bay. In order to make this decision it needs information, i.e. it needs to perceive its environment. The sensor installed in the tunnel segment in question, in particular, is responsible for providing the vehicle with the necessary information. It will present the agent, more specifically, with a picture of the traffic in that segment, by providing it with the following data:

- a list of vehicles travelling in the segment with the same direction of the vehicle under consideration, labelled GoingOutbound (GO)
- a list of vehicles travelling in the segment with the opposite direction to the vehicle under consideration, labelled GoingInbound (GI)
- the last vehicle travelling outbound in the segment, labelled LastOutbound (LO)
- the last vehicle travelling inbound in the segment, labelled LastInbound (LI)
- a list of vehicles waiting to enter the segment and which have the same direction of the vehicle under consideration, labelled QueueOutbound (QO)

- a list of vehicles waiting to enter the segment and which have the opposite direction to the vehicle under consideration, labelled QueueInbound (QI)

The behaviour of a vehicle which has to decide whether to enter a tunnel segment or wait is presented by means of the decision diagram in Figure 6.6.

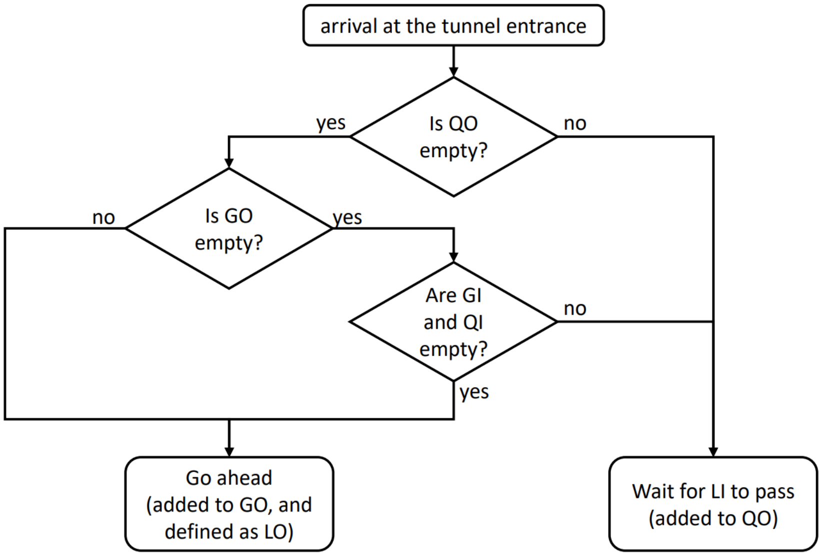

Figure 6.6 Behaviour of software agents in relation to scheduling

If there is already at least one vehicle travelling in the same direction waiting to enter (QO), according to the FIFO rule the vehicle will join the queue and will be added to QO. Otherwise the vehicle checks if there are cars with its same direction travelling in the tunnel (GO): in this case the vehicle will follow them, entering the tunnel. This choice represents the only exception to the FIFO rule (the vehicle could in fact gain priority over vehicles travelling in the opposite direction and that arrived before it in the segment, inserted in QI) and was a design choice of the author. If both lists QO and GO are empty, the agent will be the first vehicle travelling outbound, and will therefore assume leadership in coordination with vehicles travelling inbound. The lists GI and QI will thus be analysed. In case both are void, it means that the tunnel is free: the vehicle will enter the tunnel and will be added to GO and defined as LO by the sensor. If, on the contrary, at least one of the lists GI and QI contains vehicles, the FIFO rule will be followed: the vehicle will not enter the tunnel but will wait for LI to pass, and will be added to QO by the sensor.

The main advantage of the proposed scheduling strategy is the elimination of the back-and-forth trips of the vehicle backing into the conflict, which are present in the as-is mining system (left-hand side of Figure 6.5). Since the eliminated movements are non-value-added activities, the proposed approach aligns with Lean Thinking. The effects of this improvement action on the mine KPIs will be assessed thanks to the simulation project, and will be presented in Chapter 7.2.2.

6.4 Comprehensive agent-based fleet management system

Section 6.2 introduced the backbone of the solution proposed by this thesis to the fleet management problem: the agent-based system. The focus was placed on *who* is responsible for the management process, namely the software agents of the multi-agent system. Section 6.3 in the meanwhile highlighted some individual and social characteristics of the agents that define *how* these actors make the three management decisions (i.e. dispatching, routing and scheduling decisions). From the entanglement of these two aspects springs the comprehensive proposed fleet management system, an agent-based system tailor-made for underground mines. The latter is a distributed management software which supports decision-making processes in a dynamic fashion, and whose architecture is made of software agents responsible for parts of the management problem. Table 6.1, in this respect, outlines who the main agents of the agent-based system are, and which parts of the management problem they are accountable for.

Table 6.1 Agents in the multi-agent system and associated responsibilities

Agents	Dispatching	Routing	Scheduling
Vehicle agent	X	X	X
Site agent	X		

A thorough overview of the proposed agent-based system will be outlined in the remainder of this chapter. For presentation purposes, the object of study will be examined from the following perspectives: agentification (designation of agents in the system and rationale behind); agent environment; agent behaviour; interaction mechanisms.

Agentification

When conceiving an agent-based management system in which there is a problem of allocating operations to resources (as is the case in this thesis), the first stage of the design process is the agentification. With the agentification phase the problem is disentangled into decision sub-problems, that will eventually be assigned to the agents of the system. The output is thus the designation of agents, which may be permanent or temporary, and which may play different roles in the system (e.g. mediating agents, agents providing information, agents directly involved in planning, etc.).

One of the following basic perspectives can be employed to conduct the agentification phase: resource perspective; product perspective; hybrid perspective. It is worth noting that in the mining system under study the term “resource” refers to mining vehicles, while the term “product” refers to sites. In a resource perspective, resources take the lead in the planning process and allocation sub-problems: it is up to the resource agent to schedule, in dynamic or planning mode, the activities to be performed. In a product perspective, on the contrary, the leadership is taken by the product, and the sub-problems concern the identification of the resources that will be used to realize the product: it is up to the product agent to select/reserve, in dynamic or planning mode, the resources. In a hybrid perspective, finally, the two philosophies are enmeshed: allocation sub-problems can be associated to products and resources at the same time. This last perspective is the one employed in the proposed multi-agent system. On the one hand, in fact, it is the mining vehicle that dynamically schedule the activities to be performed, choosing one of the proposals received from the sites in the DynCNET protocol (Figure 6.1). On the other hand, sites have the possibility of withdrawing from the temporary agreement and dynamically choosing new, more convenient resources capable of performing the task (please refer to the behaviour of site agents in the DynCNET in Figure 6.2).

In a nutshell, by performing the agentification phase, two types of agents were determined: vehicle agents and site agents. Both are permanent agents, i.e. agents whose life cycle is linked to permanent elements of the mining system, and both are directly involved in the scheduling process, as envisaged by the DynCNET protocol. In the proposed mining system, to conclude, each site and each vehicle in the fleet will therefore be equipped with and controlled by its own software agent.

Environment of agents

The definition of agent environment is linked to the concept of information: the environment in fact represents the set of elements that the agent perceives. Given a vehicle agent or a site agent, the environment includes:

- Its virtual physical environment

  It consists of the information the agent knows, or believes to be true, about its manufacturing and logistics environment. The length of alternative routes in the routing decision is an example of such information.

- Its social environment

  It consists of the information the agent knows, or believes to be true, about other agents. A site agent, by way of example, is aware of which vehicle agents are available (i.e. not busy) at a given time.

Behaviour of agents

With regard to the dispatching part of the fleet management problem, the event- and time-driven behaviour of both vehicle agents and site agents has been defined thanks to the statecharts supporting the DynCNET protocol (please refer to Figure 6.2 in Section 6.3). The decision-making processes followed by vehicle agents in relation to the scheduling and routing parts of the problem, on the other hand, have been defined with the decision diagram (Figure 6.6) and the pseudo-code respectively (also in Section 6.3). Table 6.2 summarises the behaviour of the agents in the agent-based system, by making the overall behaviour of each agent emerge from the integration of the single behaviours adopted in relation to each of the three fleet management problems.

Table 6.2 Behaviour of software agents

Agent behaviour	Dispatching	Routing	Scheduling
Vehicle agent	Statechart (Figure 6.2)	Pseudo-code (Algorithm 1)	Decision diagram (Figure 6.6)
Site agent	Statechart (Figure 6.2)		

Interaction mechanisms

Table 6.3 outlines the main coordination mechanisms at the core of the proposed multi-agent system. For each of them, the following are specified: the medium through which the interaction takes place; the agents involved; the management sub-problem addressed with the coordination mechanism.

Table 6.3 Coordination between software agents

Coordination Mechanism	Coordination Medium	Agents involved	Management sub-problem
DynCNET protocol	Direct communication between agents	Vehicle agents and Site agents	Dispatching
FIFO rule	Virtual environment of agents	Vehicle agents	Scheduling

In relation to the dispatching part of the fleet management problem, the interaction mechanism underlying the proposed dynamic multi-agent scheduling approach is the DynCNET protocol, presented in Section 6.3. Vehicle agents and site agents therefore coordinate exploiting a protocol, which is implemented by the agents by sending and receiving messages to each other.

While as regards dispatching both the mechanism of interaction and its medium have already been presented in detail in Section 6.3, that chapter, although adequately presenting the coordination approach envisaged for scheduling, did not sufficiently emphasise the related coordination medium.

It has already been stressed, in fact, that the FIFO rule is the coordination mechanism by which access to the underground tunnels is regulated. It should be pointed out, however, that to play a fundamental role in the traffic coordination of vehicle agents is their virtual environment, which enables them to share information and coordinate their behaviour.

To decide whether to enter a tunnel segment or not (please see Section 6.3), in fact, agents monitor their virtual environment. They obtain from the sensors, in particular, the lists of vehicles in the segment in question, and on the basis of this information they decide. The decision to enter a tunnel or not, therefore, allows the agents to coordinate in traffic, in order to cope with the nature of the transport network made of one-lane bidirectional tunnel segments. The virtual environment thus serves as a flexible coordination medium: agents coordinate by adding and removing marks in the local virtual environment (e.g. when an agent enters a tunnel it signs up in the corresponding list, and when it leaves the tunnel it deletes itself from the list), and observing marks from other agents.

CHAPTER 7 RESULTS

7.1 Introduction

In this chapter, the author will first present the results found with the simulation study (Section 7.2). The focus will then be placed on the findings of an analysis conducted in parallel with the simulation project (Section 7.3).

7.2 Simulation results

The creation of a simulation model of a real underground mine was one of the key milestones of this thesis. It indeed enabled the multi-agent fleet management system proposed by the author to be simulated. Namely, the solution designed for the fleet management problem (Chapter 6) was implemented in the digital replica of the mine created with Anylogic (Chapter 5.4), in order to reproduce its behaviour and assess its management capabilities by conducting experiments in a risk-free world. The purpose of this chapter is precisely to present the reader with the findings of the simulation study. The impact of the envisaged management system on the mining key performance indicators (KPIs) will be analysed from a quantitative point of view. The KPIs values obtained with the proposed solution will also be benchmarked against those resulting from the current management system at the mine under investigation. The chapter is structured as follows: first the key performance indicators of interest will be introduced; then the results of the simulation project will be outlined.

7.2.1 Key performance indicators

The underground mining world is awash with key performance indicators on the efficiency and effectiveness of the mining logistics and production system. The KPIs employed in this thesis were selected in line with the research focus and objectives, and seek to mirror the multi-faceted nature of the problem addressed. They are mapped hereunder.

Effectiveness KPIs

- Tonnage extracted (T)
  The mine extracts four types of minerals: waste, low-grade ore, high-grade ore, super-high-grade ore. Indicating the different products as $i = 1, 2, 3, 4$, the indicator T can be defined as follows:

$$T = \sum_{i=1}^{4} T_i$$

For each of the output materials of the mine, in particular, T_i denotes the tonnes extracted and transported to surface after one working day (please remember that the conducted simulation study is focused on one business day). The indicator T thus measures the production volume of the mine, gauging the core output of the underground logistics and production system.

- Number of performed development rounds (R)
  R denotes the number of development rounds (Figure 1.1 in Chapter 1) performed by the mine in one working day. Carrying out a development cycle at an underground site, more in particular, means increasing the length of that underground tunnel and preparing the tunnel for the next cycle. R, therefore, assesses the effectiveness of the mine in terms of capacity to expand the network of underground tunnels.

- Percentage waiting time of an average site (PWT)
  The life of an underground site can be seen as the succession of the following states: waiting for the driller; being drilled; waiting for the charger; being charged; waiting for the pair scoop-truck; being unloaded; waiting for the bolter; being bolted; waiting for the cleaner; being cleaned; waiting for the driller again; and so forth. The total time that a site spends in the mine, therefore, can be broken down into two elements: a time when it waits for the resource(s) to arrive, and a time when it is processed by the resource(s). With this in mind, it is easy to grasp the meaning of the indicator PWT, which measures the percentage waiting time of an average site, i.e. the share of total time an average site spends waiting for resources. In particular, given a site i, PWT_i is defined as follow:

  $$PWT_i = \frac{WT_i}{WT_i + PT_i}$$

  WT_i, which stands for waiting time, is the time the site i spends waiting for resources to arrive, while PT_i, which stands for processing time, is the time the site i spends being processed by resources.

  To understand why PWT has been considered as a metric related to the effectiveness of the mining system, one can look at underground sites as “clients” of the logistics and production system. They require, in fact, certain activities to be performed by the resources of the mine, according to production plans: when a certain resource is requested, the site will create an order for the mine. A site that has just finished being drilled and requires, as established in the production plans, to be charged with explosives, for example, will issue an order to the mine requesting the arrival of a charger. But then, if sites are clients placing orders, an indicator measuring their waiting time (which can be seen as the time from the placement of the order, in which the arrival of a resource is requested, to the moment when the order is fulfilled, i.e. the resource has arrived) is by nature related to the mine effectiveness, as it measures the speed at which the mining system fulfils customer orders.

The effectiveness of the mining logistics and production system will be measured in the rest of the thesis by considering the three KPIs just presented. The greater the production volume, the greater the elongation of the underground tunnels, and the shorter the time required to fulfil orders from underground sites, the greater the effectiveness of the mine.

Efficiency KPIs

- Travel Efficiency (TE)

  The modelled fleet of the industrial partner is a heterogeneous group of 28 vehicles. If we denote the vehicles as $j = 1, 2, ..., 28$, and the distance travelled in one day by vehicle j of the fleet as d_j, we can define the indicator D as follows:

  $$D = \sum_{j=1}^{28} d_j$$

  D therefore measures the total distance travelled in one working day by the fleet. Remembering that the indicator T denotes the total tonnes of ore extracted in one day, we can define TE with the following formula:

  $$TE = \frac{D}{T}$$

  The metric hence indicates how many kilometres the fleet travels to obtain one tonne of mined ore, i.e. how much input must be provided to yield one unit of output.

  For the same effectiveness (T), a mine with a lower TE is more efficient than a mine with a higher TE, as it is able to achieve the same output with less input.

- Resource utilisation (U_f)

  The fleet of vehicles can be broken down into the following homogeneous groups, called families: drillers, chargers, LHDs, trucks, bolters, cleaners. Denoting the families as $f = 1, 2, ..., 6$, the utilisation (U_f) for family f is given by the arithmetic mean of the utilisation of the resources belonging to that family. The utilisation of a resource, more in detail, is given by the ratio between the busy time and the available time. The busy time refers to the time in which the resource is used by the mine, i.e. it works at a site or moves through the transport network to reach a destination. The available time of a resource, on the other hand, accounts for the total time in which that resource could potentially work: given a working day, end-of-shift breaks and restroom breaks taken by workers are excluded from the total time, giving rise to the available time. The share of available time in which the resource is not used is the idle time, which is due to factors that can either be controlled or uncontrolled by management. As will be illustrated in the section on experiments, the multi-agent management system will be able to cut out part of the idle time due to mismanagement of the company, increasing resource utilisation.

- Percentage conflict time (PCT)
  Given an average vehicle in the fleet, the PCT indicator measures the percentage conflict time, given by the ratio between the time spent conflicting with other vehicles in traffic (CT, which stands for conflict time) and the total time spent moving in the transportation network (MT, which stands for movement time).

  $$PCT = \frac{CT}{MT}$$

  Due to the nature of the transport network, which is made up of one-lane bidirectional tunnel segments, during their journey vehicles have in fact to coordinate with each other in order to resolve traffic conflicts (e.g. one car has to back into the bay to let the other pass). The travel of a vehicle within the underground mine, therefore, is given by the dichotomy between a part of the journey in which the travelled section of the transport network is clear (i.e. not occupied by other vehicles of the fleet) and a part of the journey in which the vehicle must coordinate with other vehicles in traffic. The part of the total travel time devoted to conflict resolution is the object of study of the indicator: PCT measures exactly the percentage conflict time, taking into account an average vehicle of the fleet.

  The indicator has been included in the efficiency section because the time spent by vehicles in resolving traffic conflicts can be considered as an opportunity cost: the time “wasted” by a vehicle backing into a bay to let the other one through (which represents non-value-added time) could be transformed into value-added time with better management.

Management KPIs

The six key performance indicators presented so far assess the performance of the management system in an indirect manner, i.e. by measuring the performance of the managed system (underground mine). In order not to take the conceived fleet management system only at face value, the author then decided to introduce another indicator, capable of directly assessing management performance.

- Management Quality Rate (MQR)

  This indicator aims to measure whether the company is able to send the right resources to the right working areas at the right time. This capability is synonymous with quality of the fleet management, as the latter deals with decisions on where to dispatch resources (dispatching), which route to take (routing) and how to coordinate traffic (scheduling).

  Underground working areas require the presence of mine resources at specific times, according to their status in the development cycle and production plans. A site that has just finished being drilled and that, as per production plans, needs to be loaded with explosives, by way of example, will demand the arrival of a charger. The time when a resource is requested can be defined as time of need, denoted as t_{need}. The fleet management system, on the other hand, is in charge of sending resources to working areas. We can denote the arrival time of the resource at the site with $t_{arrival}$. As one can easily guess, a management system capable of making $t_{arrival}$ coincide with t_{need} is ideal, but unrealistic. Resources will arrive at the underground working areas either early or late. Figure 7.1 shows the situations just discussed.

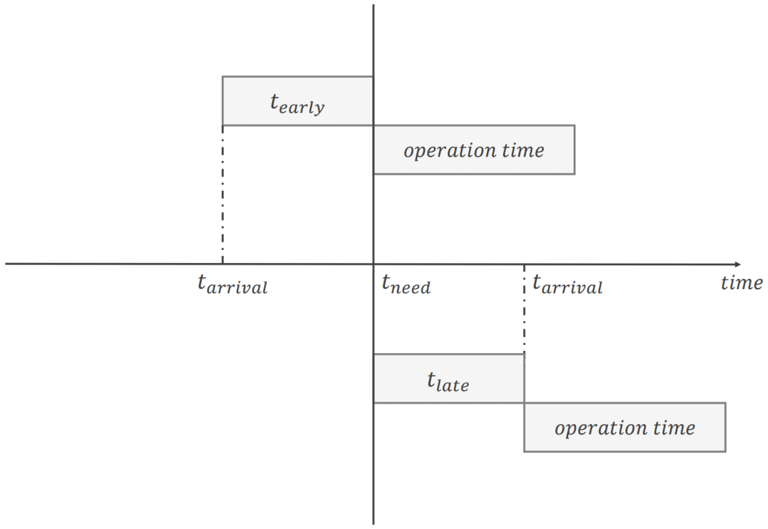

Figure 7.1 Resources arrive at underground sites either early or late

We can denote with OT the duration time of the operation, which will start once both the resource and the working area are ready. The value $|t_{arrival} - t_{need}|$ can be defined as the poor management time (PMT): the higher this value, the more the company is unable to send the right resource to the right place at the right time, i.e. the more the management system deviates from the ideal management system in which t_{need} and $t_{arrival}$ coincide. Given OT and PMT, the author defined the management quality rate as follow:

$$MQR = \frac{OT}{OT + PMT}$$

MQR ranges from 0% (very bad management) to 100% (very good management). The graph in Figure 7.2 shows the MQR value (the darker the colour, the lower the value) for different combinations of OT and PMT. Given an operation that lasts OT, the higher the PMT the worse the management. Given a PMT, vice versa, the lower the OT, the more relevant the PMT and the worse the management.

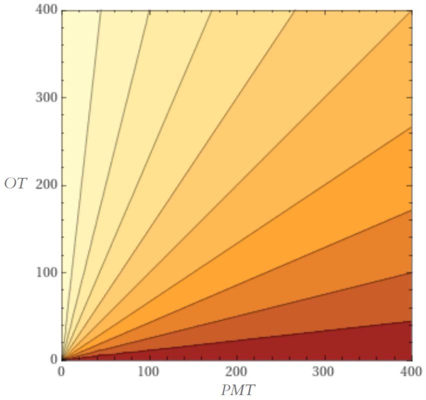

Figure 7.2 Value of MQR (the darker the colour, the lower the value)

7.2.2 Simulation experiments

The solution to the fleet management problem proposed by this thesis is a multi-agent management system in which agents enact specific dispatching, routing and scheduling strategies. Both the multi-agent system and the three management strategies have been unveiled in Chapter 6. The objective now is to quantitatively assess the performance of the proposed solution, by means of the created simulation model: the conceived management system is implemented in the digital replica created with AnyLogic, in order to study its impact on the main KPIs of the mine, outlined in the previous section.

Four experiments, in particular, are conducted: in the first three, the three strategies of dispatching, routing and scheduling are enforced in the multi-agent system individually (first only the dispatching strategy, then only the routing strategy, finally only the scheduling strategy); the last experiment, instead, deals with the comprehensive solution, in which the agents of the multi-agent system enact all three management strategies devised by the author. The experimentation strategy was defined by looking at the fleet management system of an underground mine as a system with three degrees of freedom, related to the three strategies of dispatching, routing and scheduling. Before analysing the results obtained with the final multi-agent system in which all three conceived strategies are entangled (Experiment 4), the author considered it appropriate to implement individually, one at a time, the three management approaches. Acting upon one degree of freedom at a time, in fact, will allow a better interpretation of the results achieved with the overall solution.

The results found by applying the proposed solution will also be benchmarked against the performance of the current fleet management system at the mine under review. The latter will be referred to as the “mine as-is”.

Table 7.1 summarises the foregoing characteristics of the four simulation experiments:

Table 7.1 Design of experiments

		Experiment 1	Experiment 2	Experiment 3	Experiment 4
Management system evaluated	Multi-agent system	X	X	X	X
	Dispatching strategy	X			X
	Routing strategy		X		X
	Scheduling strategy			X	X
Benchmark		Mine as-is	Mine as-is	Mine as-is	Mine as-is

Experiment 1

In the first experiment, the multi-agent system proposed in Chapter 6 has been implemented in the digital replica of the underground mine. Only the strategy designed for dispatching, however, has been enacted, without changing the approaches currently used by the company for routing and scheduling. In the experiment, therefore, the software agents of the agent-based system have brought into play the dynamic multi-agent scheduling approach defined with the DynCNET protocol.

Table 7.2 Production volumes [tons]

KPI	Planned	Mine as is	Mine to be
T_1	750	721.09	878.65
T_2	700	547.89	712.22
T_3	4400	4281.59	5012.97
T	5850	5550.57	6603.84

As depicted in Table 7.2, the proposed mining system appears to outperform the current one in terms of production volume (T). For each ore type, in fact, the average production volume given by 10 simulations of the system to-be is higher than the average production volume given by 10 simulations of the system as-is. While the system as-is lags behind the daily production plan (the achieved quantities are smaller than those planned, for each type of product), the system to-be succeeds in respecting the plans established by the mine: the volumes planned for the day are attained, and the mining system is even able to anticipate the production intended for the following day, achieving daily production quantities exceeding those planned.

Positive results were also observed in terms of expansion of the underground network: always considering 10 simulations of the two mines, the number of performed development cycles (R) skyrocketed, increasing by 21.05% over the current performance.

The PWT of an average site, at the same time, was found to fall from the 58.33% of the system as-is to the 37.98% of the system to-be, indicating a greater capacity of the mine to rapidly meet the demand from underground sites. This result is in line with expectations: the DynCNET protocol is able to allocate resources to working areas more effectively than the current dispatching approach, in which mine engineers assign at the beginning of the shift lists of destinations to be visited during the working day. An underground site becomes indeed an autonomous agent capable, for example, of dynamically taking advantage of opportunities that emerge in the mine: if a new resource agent becomes available and is closer to the site than the resource currently allocated to it, the agent will exploit the opportunity by switching resource, and will thus wait for a shorter time. The reduction of PWT, therefore, is one of the positive effects induced by the DynCNET protocol and the dynamic multi-agent scheduling approach.

To summarise what has been discovered so far, all indicators related to mine effectiveness (T, R, PWT) have been found to significantly increase with the proposed solution. The same thing, however, cannot be said about travel efficiency (TE). The total distance travelled by the fleet (D), in fact, rises by 26.76% compared to the system as-is (always considering the average of 10 simulations), and the increase in distance travelled has a more significant impact on the indicator TE than the increase in production volume (T): the mine to-be travels an average of 0.15 kilometres to obtain one tonne of extracted ore, compared to 0.13 kilometres covered in the mine as-is.

A last analysis, to conclude, must be presented in relation to resource utilisation. In the mine as-is, LHDs and trucks constitute the bottleneck of the system, i.e. the resource, or, as in this case, the group of interdependent resources (both LHD and truck must in fact be present for unloading to take place) that is most stressed. The utilisation of these two types of resources is 100%, i.e. they are busy for the whole available time: they are therefore the production assets whose maximum capacity is insufficient in relation to the workload resulting from the mine production plans. The state of affairs does not change when implementing the management system envisaged in Experiment 1: LHDs and trucks continue to be the bottleneck of the system, with 100% utilisation. With the solution proposed in Experiment 1, however, the bottleneck of the system is able to cope with the workload imposed by production plans. Utilisation is 100% only because in the created simulation model, once the planned work is completed, mining resources anticipate the work scheduled for the next day.

What improves with the proposed solution, moreover, is the utilisation of the other resources of the mine. Excluding LHDs and trucks from the vehicle fleet and considering the remaining group, the average utilisation turns out to be 55.11% in the system as-is (again considering 10 simulations). The non-utilisation of a resource can be interpreted as the composition of two causes. The first is the performance of the bottleneck of the system. The utilisation rate of a non-bottleneck resource, in fact, is not determined by its potential capacity but by the performance of the bottleneck: LHDs and trucks clog the “production flow” having repercussions on all the other vehicles, which have to wait for the unloading activity to be completed before they can perform their operations (if, for example, LHDs and trucks are able to process 2 sites per hour, bolters, despite having a capacity of 10 sites per hour, will work 2 sites per hour, and will be unused for most of their time). The second cause is the corporate mismanagement, i.e. the poor performance of the adopted management system; in the system as-is, in fact, a resource may be unused because it is dispatched to a work zone where it must wait for the end of ongoing operations, while there may be other sites ready to be processed (there is therefore no problem in terms of bottleneck). In the system as-is, the two causes account on average for 60.64% and 39.36% of the average non-utilisation time respectively. In the system to-be, on the other hand, the management inefficiency (second cause) has been eliminated, and the non-utilization of non-bottleneck resources is explained solely and exclusively by the performance of the bottleneck of the system (first cause). In the DynCNET protocol, in fact, site agents participate in a call-for-proposal of a resource only if they are ready to receive that resource: resources therefore do not have to wait once arrived at destination, and the non-utilisation linked to management inefficiency disappears. The average utilisation of a non-bottleneck resource, in conclusion, rises from 55.11% of the system as-is to 68.22% of the system to-be.

Experiment 2

In the second experiment, the dispatching and scheduling approaches currently employed by the mine under investigation were preserved, while the designed routing strategy was implemented. That is, when several alternative routes exist to reach the same destination point, the agents of the multi-agent system choose considering both the distance factor (as done in the as-is system) and the traffic factor (this is possible thanks to the characteristics of the proposed cyber-physical system, such as, for example, the presence of sensors). We therefore move from a system in which the chosen route is always the shortest one, no matter what, to one in which agents could also opt for the longest path in case the shortest one is too congested.

Table 7.3 90% confidence intervals for D, PCT and PWT

90% confidence intervals	Mine as is		Mine to be	
	Average	Half width	Average	Half width
D [km]	748.60	39.80	751.68	41.01
PCT	5.82%	2.28%	4.66%	2.48%
PWT	58.33%	5.87%	59.92%	5.94%

Table 7.3 outlines the 90% confidence intervals of some of the KPIs affected by the adopted strategy. Considering the centres of the confidence intervals (i.e. the averages), the total distance covered by the fleet after one working day (D) slightly increases, while PCT slightly decreases. The deviations, albeit minimal, are in line with expectations: by opting for longer routes to avoid traffic on the shortest path(s), vehicles cover a greater distance and spend less travelling time in conflict with others. The strategy adopted, however, does not seem to have much impact on the waiting time of underground sites (PWT): with respect to the mine as-is, vehicles arrive at their destination more or less always with the same "punctuality". The values of the other KPIs, moreover, do not deviate significantly from those of the system as is. In light of the foregoing, the proposed routing strategy does not seem to have a significant impact on the performance of the mining system. The experiment findings, however, might be biased by a limitation of the simulation model, as will be discussed in Section 8.3.

Experiment 3

Similarly to the first two experiments, in the third experience two of the three management strategies adopted by the mine were left unchanged, in order to individually study the effects brought about by the variation of the third strategy (in this case, the scheduling one). The key objective of the designed scheduling approach, in which sensors, agents and on-board navigation systems coordinate vehicle traffic, was to increase the efficiency of the mining system, by reducing the distance travelled by the fleet within the mine. The results obtained with the simulation experiments in AnyLogic highlight the capability of the strategy to achieve the intended goal: the overall distance travelled by the fleet in a working day (D) decreases by 38.12% compared to the current performance. The increase in mine efficiency, moreover, does not entail any sacrifice in terms of "service level": the number of performed development cycles (R) considering 10 simulations turns out to be 8.98, almost equal to the 9.46 cycles carried out in the system as-is, while the total amount of material brought to the surface (T) increases by 2.49% when considering the entire production mix. As a result TE, able to account for any trade-offs between efficiency (D) and effectiveness (T), decreases from 0.13 to 0.08 kilometres travelled to obtain one tonne of extracted ore.

Experiment 4

At this point of the experimentation phase, thanks to the propaedeutic studies carried out with experiments 1, 2 and 3, we are aware that:

1. The dispatching strategy proposed with the DynCNET protocol and the dynamic multi-agent scheduling gives rise to a more effective mine than the current one. All KPIs, with the exception of TE, turned out to be improved, and moreover with significant magnitudes. The improvement of the fleet management system is radical.
2. The envisaged routing strategy does not seem to significantly affect the performance of the current mine. This may be due, however, to limitations of the created simulation model, as will be explained in Section 8.3.
3. The scheduling strategy is able to reduce the total distance travelled by the fleet, without jeopardising other performances. The improvement can be considered incremental.

With this in mind, we can now analyse the results obtained by implementing the overall solution proposed by this thesis, i.e. a multi-agent fleet management system which enforces all three of the designed dispatching, routing and scheduling strategies. The dashboard in Figure 7.4, in this respect, depicts the findings[1]. KPIs figures refer to the average values obtained with 10 simulations in AnyLogic. The mine to-be can also be benchmarked against the mine as-is thanks to Figure 7.3, which portrays the KPIs values resulting from 10 simulations of the current mine.

The results obtained in Experiment 4, as expected, are very similar to those observed in Experiment 1. The dispatching strategy, in fact, is the one that steers the performance of the mine. By combining it with the scheduling strategy, moreover, it is possible to further improve the management system, reducing the distance travelled by the fleet.

---

[1] Note: Dashboard items *average equipment*, *average scoop*, *average truck* and *average working area* refer to average resources and average working areas. The value of a KPI presented for an average truck, by way of example, is obtained by arithmetically averaging the KPI values of all the trucks in the fleet. The same reasoning applies to other items wherein the term “average” appears.

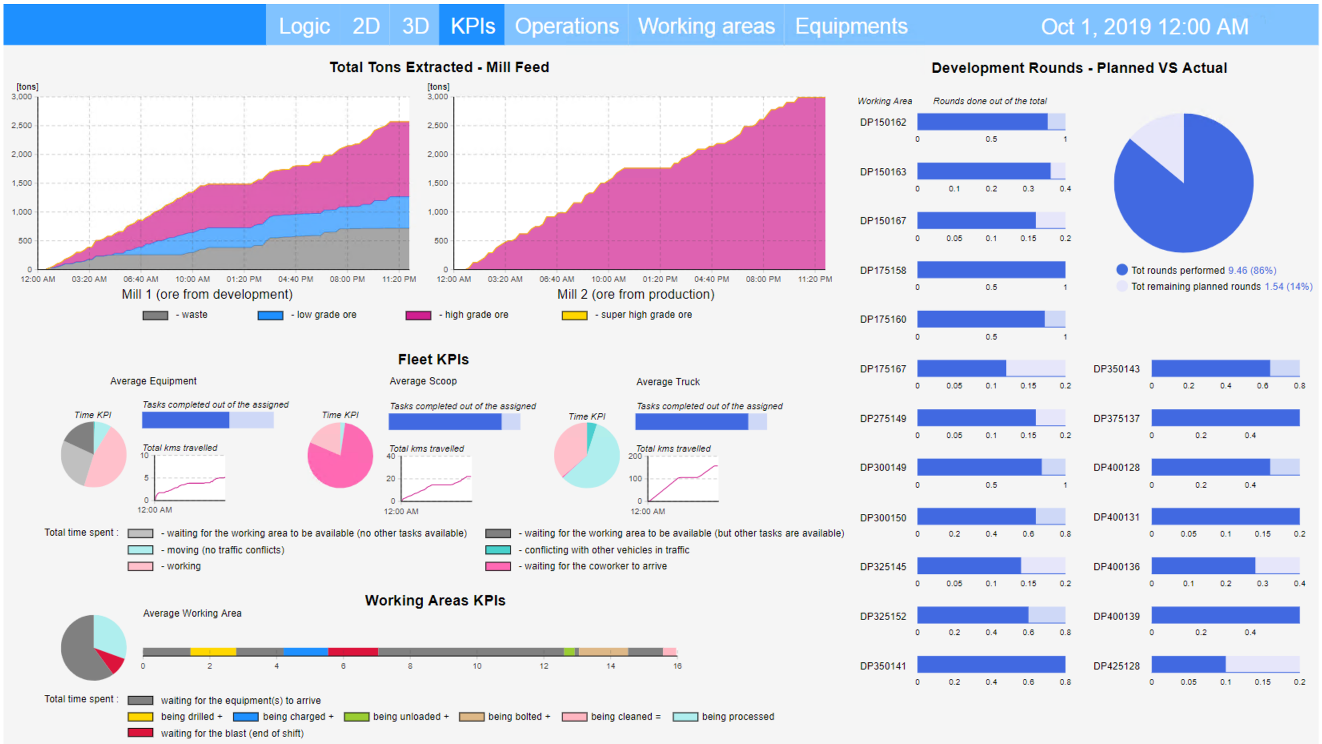

Figure 7.3 Average value of KPIs from 10 simulations in AnyLogic of the mine as-is

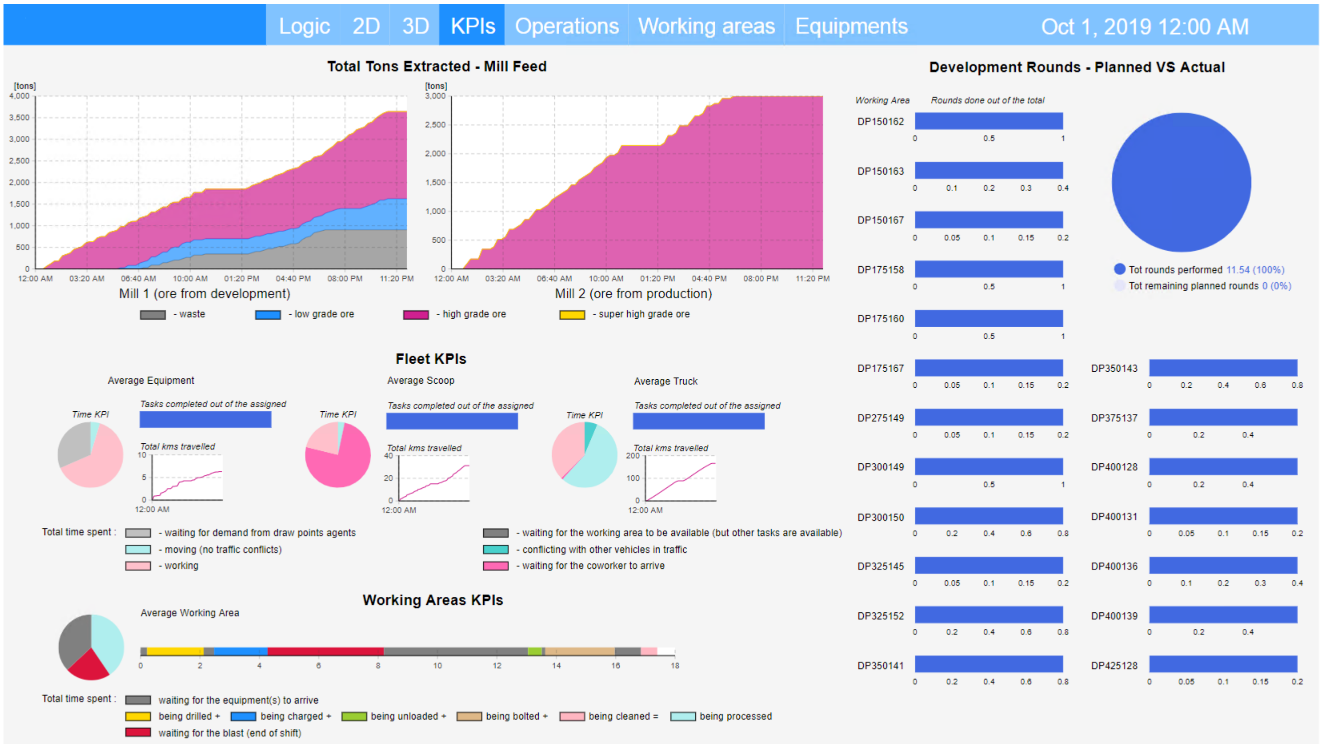

Figure 7.4 Average value of KPIs from 10 simulations in AnyLogic of the mine to-be

The cross comparison of figures 7.3 and 7.4 is evidence that the proposed agent-based fleet management system outperforms the current management approach adopted by the industrial partner.

At the end of the working day under review, the total amount of ore produced and transported to the surface (T) is 19.63% higher than the value of the mine as is. In relation to this KPI the dashboard also specifies: the composition of the production mix, indicating the extracted quantities for each ore type (waste, low-grade ore, high-grade ore, super-high grade ore); the distribution of production volumes over the day; the surface mill to which the extracted ore is transported.

The right-hand side of the dashboard then outlines the number of executed development cycles (R). We note that while the current mine is not able to keep up with production plans, achieving on average only 86% of the 11 planned cycles, the mine to-be not only manages to stay on schedule, but is even able to anticipate the activities planned for the next day, performing an average of 11.54 development cycles (against the 11 planned for the day under review).

The designed management system also yields better performance in terms of percentage waiting time (PWT) of underground sites (please refer to the heading “Working Areas KPIs”). As illustrated in the dashboards, in fact, the percentage waiting time for an average site drops from 58.33% of the current mine to 36.95% of the mine managed with the agent-based system. The logistics and production system is therefore able to fulfil demand from working areas more promptly. The bar chart under the section “Working Areas KPIs” also untangle the total waiting time into the single waiting times for individual resources: the fact that the waiting time for the LHD-truck pair is the longest, in both the as-is and to-be systems, is consistent with the nature of this resource group, which represents the bottleneck of the mining system.

As regards mine efficiency, under the heading “Fleet KPIs” it is possible to observe how resources spend their time inside the mine. Results are presented for an average equipment (i.e. an average vehicle of the group consisting of drillers, chargers, bolters and cleaners), an average scoop and an average truck. The total time, moreover, refers to nominal operating conditions, i.e. it excludes shift-change breaks. The utilisation of trucks and LHDs is 100%. In the mine as-is, in fact, LHDs and trucks constitute the bottleneck of the system, i.e. the resource, or, as in this case, the group of interdependent resources (both LHD and truck must in fact be present for unloading to take place) that is most stressed. They are the production assets whose maximum capacity is insufficient in relation to the workload resulting from the mine production plans. It should be noted that the bottleneck of the logistics and production system does not "shift" and remains the same also in the mine to-be. Thanks to the proposed management system, however, the performance of the bottleneck manages to meet the load imposed by production plans: the utilisation of trucks and LHDs continues to be 100% only because in the created simulation model, once the targets of the day have been reached, resources anticipate the activities scheduled for the following day.

A separate discussion needs to be made for the average equipment. In the system as-is, resources are provided with lists of destinations to visit during the day. Having reached a destination, however, it is possible that the site in question is still being processed by the previous equipment in the development cycle. Resources in this case will have to wait, and the waiting time will be idle time. Part of the idle time, however, may be an opportunity cost, as the resource may go to other sites already available to accommodate it, turning the idle time into busy time. Figure 7.3 shows that the utilisation of an average equipment is 55.11%, and that the 39.36% of the idle time is an opportunity cost. In the mine to-be, on the other hand, the choice of the next destination is made dynamically using the DynCNET protocol. Once the current task is completed, the equipment notifies all site agents of its availability. If no one is interested in its services at that moment, however, the resource will remain idle, until there is at least one site that asks to be treated by taking part in the protocol. Only site agents who are available to accommodate a resource will bid on it; once at destination, therefore, the vehicle is sure to be able to start its processing without having to wait for the end of the previous operations. For this reason, the opportunity cost in the mine to-be is zero: by eliminating the share of idle time linked to opportunity cost, the overall idle time drops, and the utilisation of the average equipment rises to 68.45%. The findings are in line with those already observed in Experiment 1, wherein the phenomenon was explained from another perspective: the non-utilisation of a non-bottleneck resource was seen as the composition of two causes, i.e. the performance of the bottleneck, and the mismanagement of the company. The interested reader may refer to the discussion of Experiment 1.

As already explained above, the conceived mining system is able to achieve a greater output than the current one, both in terms of production volume (T) and expansion of the underground network (R). The price to be paid, logically, is the provision of a greater input, defined in terms of distance travelled by the fleet (D). The kilometres covered by the fleet in one day, in fact, go from 748.6 km in the mine as-is to 827.2 km in the mine to-be. It should be noted, however, that the value is lower than the one obtained in Experiment 1 (948.96 km): in Experiment 4, in fact, the dispatching strategy merges with the scheduling strategy, which, as found in Experiment 3, allows to reduce the distance covered by the fleet, ceteris paribus. What is important to consider about the distance travelled, however, is the value of TE, which captures both the variations obtained in terms of efficiency (D) and those obtained in terms of effectiveness (T). The 0.13 kilometres travelled by the current fleet to yield one tonne of ore become 0.12 kilometres per tonne in the system to-be.

The "as-is" and "to-be" fleet management systems were finally benchmarked in terms of MQR, which was computed for each operation in the development cycle.

Table 7.4 MQR values

MQR	Drilling	Charging	Unloading	Bolting	Cleaning
Mine as is	46.33%	38.85%	7.06%	53.51%	22.85%
Mine to be	89.39%	84.14%	8.26%	95.71%	38.65%

Bearing in mind that the value of MQR can range from 0% (bad management) to 100% (excellent management), the results indicate that the multi-agent management system outperforms the one currently employed by the industrial partner, as it is more capable of sending the right resources to the right place at the right time (please refer to the discussion on the meaning of MQR presented in Section 7.2.1).

7.3 Qualitative analysis

In the previous section the envisaged solution to the fleet management problem was studied from a quantitative point of view, assessing its impact on the three sets of KPIs, capable of monitoring the effectiveness, efficiency, and management capabilities of the underground mine respectively. The designed multi-agent management system, moreover, offers the company some non-quantifiable benefits that are not directly reflected in the value of the KPIs, but which are equally valuable. The latter are the focus of this section.

Production planning and operations scheduling in an underground mine are major challenges, owing, inter alia, to the numerous interdependencies that exist between the various mining activities. Figure 7.5 portrays such interdependencies, and emphasises their nature following the classifications proposed by [53] and [54]. As depicted in the figure, the activities of the development cycle are linked by a *sequential interdependence*: each operation cannot start until the previous one has been completed. These constraints must be fulfilled, evidently, at each underground site. Sites, meanwhile, compete with each other for access to mining resources (*pooled interdependence/organizational resources sharing*). This is true, in particular, for each operation of the development cycle for which a (group of) equipment is required: drilling activities at different sites will contend for drillers, charging activities will contend for chargers, and so on. Finally, there is a *simultaneity interdependence* between the blasting operations, which must occur in the same time windows. The explosion can in fact only take place at the end of the work shift, when all the workers have returned to the mine surface and are safe.

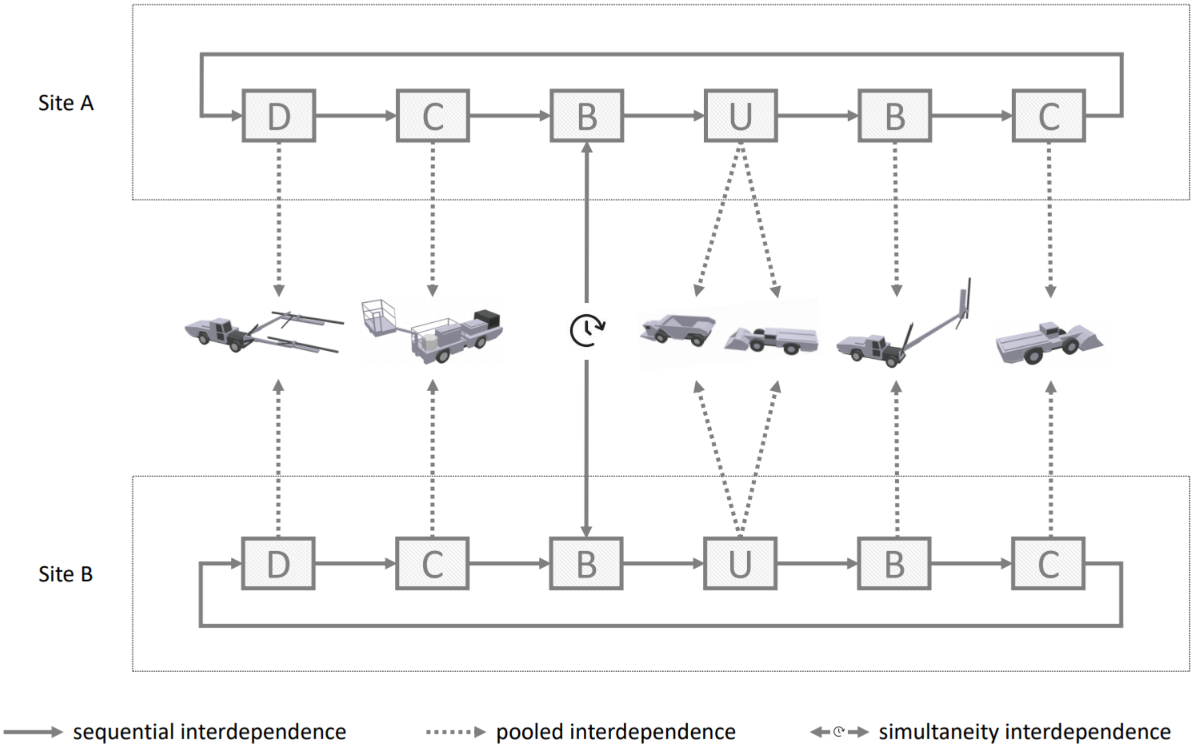

Figure 7.5 Interdependencies among mining activities (D:drilling - C:charging - B:blasting - U:unloading - B:bolting - C:cleaning)

In order to pursuit the objectives of efficiency and effectiveness, all the presented interdependencies have to be managed with some regulation and coordination mechanisms. Following the definition suggested by [54], we know in fact that coordination is the "process of managing dependencies among activities". The coordination and control mechanism set up with the proposed multi-agent fleet management system will be outlined using the classification scheme presented by [5] (Figure 7.6), which deals with coordination and control of manufacturing activities in distributed and agent-based manufacturing systems. Against this background, it is worth emphasising that in the mine under review, at the beginning of a working day, engineers provide resources with lists of work to be performed during the day. We are therefore in the presence of a *third party coordination by plan* (quadrant I-4 in Figure 7.6): engineers play the role of superiors, resources that of centres to be coordinated, and predefined plans are established to coordinate interdependent activities a priori and in a coercive way. In the designed multi-agent management system, on the other hand, vehicle agents and site agents are fully autonomous, interact with each other with the DynCNET protocol, and schedule operations dynamically throughout the working day. The resulting coordination and control mechanism, therefore, can be classified as *coordination by mutual adjustment during activities execution* (quadrant I-3).

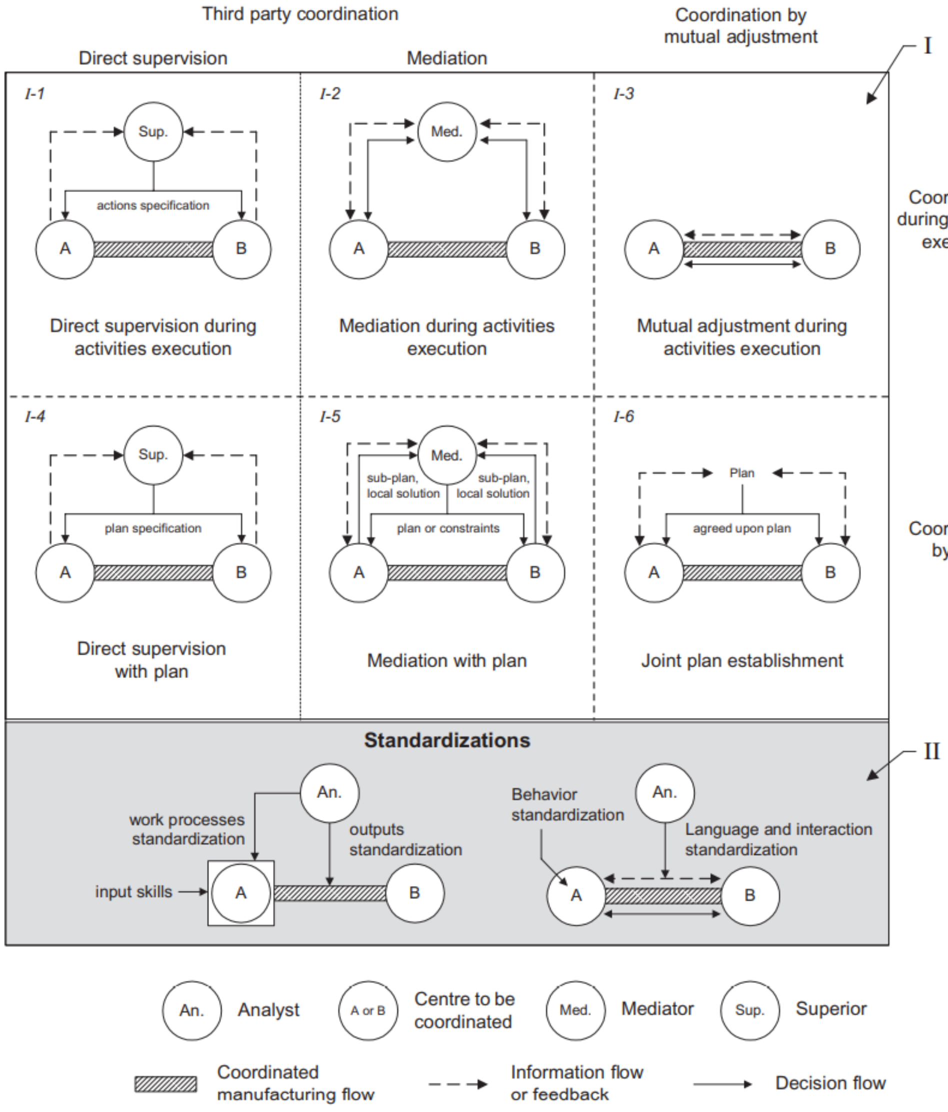

Figure 7.6 Classification scheme of coordination and control by [5]

In the light of the foregoing, we can say that: (i) the resources of the envisioned mining system are provided with much greater autonomy than at present; (ii) there is a shift from passive coordination by plans to dynamic scheduling and active coordination during activities execution.

Taking a broader perspective, organisations, in order to manage the interdependencies between production activities (Figure 7.5), put in place organisational mechanisms which shape the so-called *control architectures*, i.e. structures for decision-making and information exchange between the various elements of an organisation and/or the various equipment of a production system. [55], in this respect, defined control architectures as "patterns of decision-making and communication among a set of actors who perform tasks in order to achieve goals". [5], more in detail, propounded a list of reference control architectures, of which the following are part: centralized control architecture; proper hierarchical control architecture; modified hierarchical control architecture; heterarchical control architecture; quasi-heterarchical control architecture. The control architecture put in place by the mine under study can be considered *modified hierarchical* (Figure 7.7). At the highest tier of the hierarchy are the engineers, who create the work plans (dispatching lists) for the working day and assign them to the workers. At the subsequent level are the mine workers, each responsible for a vehicle: although the working areas to be visited during the day are determined by the lists received from the upper level, these control components have a certain degree of autonomy in management and control, as they can decide which route to take to reach the (already established) destination, and how to coordinate with other vehicles in the traffic. In other words, the first level of the control architecture is in charge of the dispatching decision-making process, while the second layer takes care of the routing and scheduling ones.

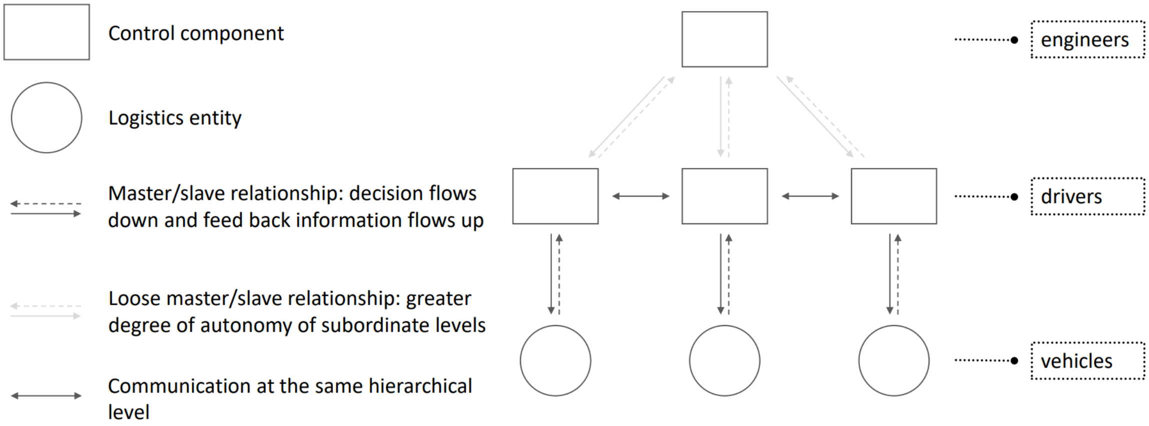

Figure 7.7 Modified hierarchical control architecture

The control architecture envisaged with the multi-agent system, on the other hand, can be deemed *heterarchical* (Figure 7.8), since management and control responsibilities are locally distributed to autonomous control units that interact with each other without superior/subordinate relationships [5]. Site agents and vehicle agents structuring the agent-based system are in fact the sole control components of the architecture, and are indeed endowed with maximum autonomy. Vehicle agents, for example, choose autonomously which working area to visit from time to time (i.e. dynamically), which route to take, and how to coordinate in traffic. In other words, all three decision-making processes of dispatching, routing and scheduling are the responsibility of the individual actors of the logistics and production system.

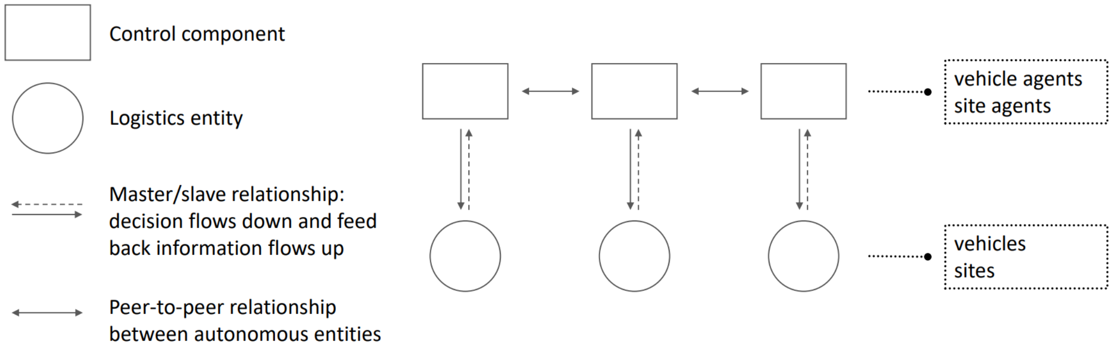

Figure 7.8 Heterarchical control architecture

In the transition from the mine as-is to the mine to-be, therefore, we move from a *modified hierarchical* control architecture to a *heterarchical* control architecture, and resources in the distributed mining system gain a much greater degree of autonomy than at present.

In the light of all that has been said so far in terms of both coordination mechanisms and control architectures, the proposed multi-agent system, in contrast to the current mining system, exhibits two properties: autonomy of the individual resources in the mining system; dynamic operations scheduling. There is thus a shift from workers executing pre-established plans to autonomous workers creating plans; there is a shift from operations scheduling to dynamic operations scheduling.

Thanks to these two characteristics, the underground mine would acquire the following qualities:

- Agility
  This thesis puts forward a distributed vision of the mining organisation composed of autonomous units, thus promoting agility of the mining logistics and production system. One of the ways to foster organisational agility, in fact, is to give more autonomy to the different elements of the organisation.
  Agility can be seen as the property of a complex system capable of rapidly exploiting information from its volatile (unstable and unpredictable) environment and to self-adjust in order to effectively and efficiently achieve its objectives. Being agile therefore means being reactive, resilient, flexible and adaptable. All these qualities are of paramount importance in the world of underground mines, due to the high volatility inherent in these settings. Underground mines, for example, are subject to a variety of unexpected disruptions (resource-related, job-related etc.) which can change system status at any

time and affect its performance.

Delving into details, in the current mine dynamism is tackled in a reactive and centralised manner: real-time information continually forces engineers to revise and adapt pre-established schedules. With the multi-agent system, on the other hand, dynamism is part of the nature of the system itself (let's consider, for example, the dynamic multi-agent scheduling approach), and is managed in a proactive, decentralised and distributed manner. Literature, in this respect, is full of evidence that distributed management systems, such as the one presented in this thesis, are more suitable for managing the dynamics of complex processes than centralised systems.

- Self-organisation

  From the autonomy inherent in the multi-agent system stems its capacity to self-organise. An example of such a capability is provided by the dynamic operations scheduling undertaken by vehicle agents and site agents. Self-organisation takes the shape, for instance, of self-coordination (i.e. the capability to control one's own actions in function of others): an example is provided by the routing decision taken by vehicle agents, that choose which route to take in function of the traffic.

It is also worth noting that in the conceived solution the autonomy of the individual units of the system is not in trade-off with the unit of action of the system. Generally speaking, in fact, autonomous behaviour, if entirely self-determined, can endanger the pursuit of the overarching objective of the organisation and promote anarchies. The problem of aligning the autonomy of individuals to the common purpose of the whole organisation is well known both in the literature of distributed systems (a reference is [5]) and in that of underground mining (a reference is [6]). A strategy suggested in both contexts is to implement a system of dual goals for the individual units of the organisation: on the one hand there are local objectives, which leads each unit to act autonomously according to what is best for itself; on the other hand there are objectives linked to the overall performance of the system, which ensure a certain global vision. The proposed multi-agent system embraces this dual perspective: in the DynCNET protocol, for example, vehicle agents choose the next destination to visit considering on the one hand the distance to be covered (local goal), and on the other hand the priorities set by production plans (global goal). In this way, the autonomy of the individual actors in the mining system leads to organisational agility and self-organisation without jeopardising the unity of action.

To summarise what has been said throughout this section, the multi-agent system proposed for the fleet management problem, in addition to the quantitative advantages reflected in the value of the KPIs (Section 7.2.2), is capable of giving rise to a logistics and production system that is agile and capable of self-organisation. The significance of these properties is paramount in the world of underground mining. The literature is in fact awash with references to the high uncertainty, turbulence and unpredictability in underground mines. Shaping a mining system with the aforementioned qualities is therefore a major milestone.

CHAPTER 8 DISCUSSION

8.1 Answering the research question

As this research project draws to a close, a summary of the major achievements can be put forward, in order to highlight how the research objectives have been progressively met and how the research question from which it all began has been answered.

The first milestone of the project, in this respect, was the design of the solution to the fleet management problem, grounded on agent-based systems, Industry 4.0 and dynamic multi-agent scheduling. This enabled the first research sub-objective (RSO1) to be achieved. In a parallel process, the digital replica of a real underground mine (industrial partner) was built with AnyLogic. The conducted simulation project made it possible to reproduce the behaviour and assess the performance of the proposed fleet management system (RSO2). The simulation model also allowed to compare the designed agent-based system with the management strategy currently adopted by the industrial partner (RSO3). The last research sub-objective (RSO4) was finally achieved by qualitatively comparing the proposed solution with the main ones presented in the underground mining literature (Chapter 2) and with the one enforced by the industrial partner (Chapter 7.3).

In the light of the findings of the simulation study, the designed agent-based fleet management system appears to outperform the one currently employed by the company under investigation. The underground mine, in fact, proves to be able to achieve higher production volumes, to have a greater capacity for expansion of the underground network, and to fulfil the demand from underground working areas more promptly. All these factors are synonymous with enhanced effectiveness. At the same time, bottleneck performance improves, resources attain higher utilisation, and the fleet turns out to be able to yield one tonne of extracted ore by travelling fewer kilometres. All these factors are synonymous with enhanced efficiency. By adopting the solution proposed by this thesis, in a nutshell, the mine would benefit from a more efficient and effective logistics and production system. In view of the foregoing, it seems rational to give an affirmative answer to the research question that initiated this thesis (illustrated below), thus coming full circle.

RQ: Can agent-based systems and dynamic multi-agent scheduling solve the fleet management problem in an underground mine, leading to an efficient and effective logistics and production system?

The remainder of this chapter is structured as follows: first, the author will present a critical reflection on the results of the simulation project; secondly, the limitations of the study will be highlighted; finally, potential avenues for future research will be outlined.

8.2 Critical reflection

In Chapter 7 ("Results") the agent-based fleet management system proposed by this thesis was evaluated from two viewpoints. On the one hand, thanks to the simulation project, it was possible to estimate its impact on the KPIs of the industrial partner. On the other hand, the spotlight was placed on intangible properties (agility and self-organisation) which are not directly reflected in the value of the KPIs, but which are equally meaningful.

With reference to the first perspective, the results of the simulation study indicated that the main KPIs of the underground mine, related on the one hand to the efficiency and on the other to the effectiveness of the logistics and production system, improve by adopting the proposed solution. The key objective of simulation experiments, however, was not so much to find out whether with the designed agent-based system the performance of the mining system improved or worsened, but rather to assess the extent of the improvement. Adopting the solution suggested by this thesis, in fact, would mean investing in automation and Industry 4.0, taking the first steps towards the "smart mine", in pursuit of the benefits and value proposition that this paradigm claims. In return for the investment, therefore, the improved performance of the mine was something to be expected even without conducting the simulation experiments. The strength of this study, instead, is rather to assess the magnitude of improvements, and thus help the company to evaluate the return on investment. While in the past industrial revolutions, in fact, the choice "invest vs. not invest" was quite straightforward because the extent of both the investment and the returns was quantifiable, today most companies are sceptical about investing in Industry 4.0 precisely because it is difficult to quantify the benefits that will be gained, and to foresee, thus, the returns on the investment. What is the advantage, for example, of possessing real-time information on the state of the logistics and production system and making decisions based on it? In the agent-based system described in this research, for example, software agents have global visibility of the real time status of the entire underground mine, and, based on this knowledge, make decisions (an example of decision-making is the dynamic scheduling of operations). Although it is obvious that such a system brings benefits and improves the state of affairs, the hurdle to investment is precisely to quantify the extent of these benefits. In this thesis, such an obstacle was overcome with the aid of simulation. The value proposition offered by this research project, in a nutshell, is precisely to quantify the benefits of investing in Industry 4.0, and thus justify the investment: thanks to the simulation project and the digital replica created with AnyLogic, we have in fact managed to capture the values of the increase in production volumes, service level, productivity and resource utilisation, among others. This, according to the author, is the main exploit of the thesis.

8.3 Limitations

This project is limited to a specific industrial case and the objective of the experiments is confined to the analysis of the performance of the proposed solution for such an industrial case.

The other main limitations of this study are related to either the simulation model or the agent-based system.

First, the digital replica of the mine does not model the event-based rescheduling activities undertaken by mine engineers. When a vehicle reaches an underground site and there are still operations in progress, in fact, in the created simulation model the resource will wait for the end of the latter, and then begin its activity. In the real world, on the other hand, if the waiting time is not negligible, the worker will report the matter to the engineers, who, based on production plans and the current state of the mine, will decide on the countermeasures to be taken. The vehicle could thus potentially be dispatched to another site. Modelling the engineers' decision-making process is a non-trivial endeavour that was not carried out in the following research project, but leaves room for future research. Albeit the outcomes of the simulation of the mine to-be (i.e. the one in which the agent-based system is implemented) are not tainted in any manner by the foregoing modelling deficiency, the results of the simulation of the mine as-is could be biased. For this reason, in order to assess the impact of the model limitation on the results of the simulation of the mine as-is, the author analysed the number of occurrences of the (potential) problem, i.e. the number of times during simulation that resources arrive at destination, there are still operations in progress, and it is necessary to wait a non-negligible amount of time before processing can begin. It was found that, out of the total number of dispatches (i.e. the total number of times resources arrive at destination), only 5.85% of the time does this (potential) problem occur. We can therefore conclude that the model limitation does not significantly impair the results presented for the system as-is, also considering the fact that the 5.85% refers to problems that are only potential. It is not necessarily the case, in fact, that, given a resource that has to wait a not negligible amount of time, engineers in the real world decide to dispatch the resource to another destination. However, despite the fact that the distortion is not disruptive, if one wishes to read the simulation results (Section 7.2.2) taking into account the noise introduced by the model limit, the following guide can be referred to: the KPIs figures for the mine to be, as already explained, are not distorted; for the mine as is, instead, the values of R, TE and U_f may slightly underestimate the actual figures, the value of T is not affected by the model limitation, and the effects of the model limit on PWT, PCT and MQR cannot be defined a priori.

The second limitation of the simulation model concerns the underground transport network. The layout of the partner mine, in fact, does not feature many alternative roads, which implies a low number of occurrences of the routing decision (choice of which road to take to reach a destination). Implementing the conceived routing strategy in the agent-based system (experiment 2) consequently led to results which, although in line with expectations, did not differ radically from those of the mine as-is. In order to evaluate the proposed routing strategy more thoroughly, therefore, in future works the layout of the digital replica of the mine will be expanded so as to increase the number of alternative routes and, as a consequence, the number of times the agents of the agent-based system are faced with the routing decision.

The last limitation of the study concerns the designed agent-based system. The key objective of this thesis was to evaluate the potential of agent-based systems and dynamic multi-agent scheduling procedures in solving the fleet management problem in underground mines. Consequently, while for the dispatching part of the problem a fancy strategy such as the DynCNET protocol has been proposed, the strategies devised for the routing and scheduling aspects of the problem are quite simple, and have room for improvement. In future works, more elaborate routing and scheduling strategies may be designed.

8.4 Future research

Although the aforementioned limitations restrict the scope of this project, they also represent valuable research opportunities. The main research direction should head towards exploring the performance of the agent-based fleet management system in mining contexts other than the specific industrial case treated in this research project. Another avenue of future research could focus on improving the proposed agent-based system. Vehicle agents and site agents, for example, could become intelligent agents, i.e. agents able to learn from their past and dynamically adapt their future behaviour: dispatching, routing and scheduling decisions would then be made based on both past experience and current inputs. Last but not least, the most ambitious project would be to evolve the digital replica of the mine into a digital twin. This would mean introducing a two-way exchange of information between the physical world and the simulation model created with AnyLogic. Simulations would leverage real-time data to mirror the physical world in the virtual model, and simulation results would be used by actors in the logistics and production system to make real-time decisions (e.g. those concerning dispatching, routing and scheduling). Put another way, simulation experiments would be executed in real time in order to predict the behaviour of the mining system in different scenarios, and make consequent decisions on-the-fly. If this avenue of research were to be pursued, the employment mode of simulation would comply with that envisaged by the Industry 4.0 philosophy: on its journey towards the mine 4.0, therefore, the company would go even further than the first steps already taken with this research project.

CHAPTER 9 CONCLUSION AND RECOMMENDATIONS

In this research project, an agent-based system and a dynamic multi-agent scheduling strategy tailor-made for fleet management in underground mines have been designed.

In order to reproduce the behaviour and evaluate the performance of the proposed agent-based fleet management system, an agent-based simulation model of a Canadian underground gold mine has been built using AnyLogic.

The results of the simulation study have indicated that, by adopting the solution proposed by this thesis, the mining system would be able to achieve higher production volumes, a greater capacity for expansion of the underground network, and a better responsiveness to demand from underground working areas than at present. This is synonymous with increased effectiveness of the logistics and production system. At the same time, the performance of the bottleneck of the system would improve, resources would reach higher utilisation, and the mining system would be able to yield a tonne of extracted ore by travelling fewer kilometres. This is synonymous with increased efficiency of the logistics and production system. Last but not least, with the agent-based fleet management system the mine would benefit from enhanced agility and self-organisation.

The core achievement of the research project, however, is not so much to demonstrate improvements in the performance of the mining system, but to quantify the extent of such improvements, by capturing the precise variations in the mining KPIs. Indeed, adopting the proposed solution would mean investing in Industry 4.0, in pursuit of the benefits envisaged by the paradigm; it seems rational, therefore, to foresee the improvement of the state of affairs even without the conducted simulation project. The main exploit of this thesis, instead, is precisely to quantify the magnitude of the achievable benefits, and to gain an insight into the returns on investment in Industry 4.0.

On the path towards the mine 4.0, in order to go even further than the first steps taken with this thesis, in future research the digital replica created in AnyLogic could be transformed into a digital twin, introducing a two-way exchange of information between the physical and virtual worlds. In this way, players in the logistics and production system would be able to make real-time decisions by conducting on-the-fly simulations, as envisaged by Industry 4.0.

REFERENCES

[1] L. T. Blessing and A. Chakrabarti, *DRM: A design reseach methodology.* Springer, 2009.

[2] A. Borshchev, Ed., *The Big Book of Simulation Modeling: Multimethod Modeling with Anylogic 6.* Chicago, IL, USA: AnyLogic North America, 2013.

[3] K. E. Drexler, *Radical Abundance: How a Revolution in Nanotechnology Will Change Civilization.* Public Affairs, 2013.

[4] R. G. Sargent, "An assessment procedure and a set of criteria for use in the evaluation of computerized models and computer-based modelling tools," Syracuse univ NY dept of industrial engineering and operations research, Tech. Rep., 1981.

[5] J.-M. Frayret, S. D'Amours, and B. Montreuil, "Coordination and control in distributed and agent-based manufacturing systems," *Production Planning & Control*, vol. 15, no. 1, pp. 42–54, 2004.

[6] M. Gamache, R. Grimard, and P. Cohen, "A shortest-path algorithm for solving the fleet management problem in underground mines," *European Journal of Operational Research*, vol. 166, no. 2, pp. 497–506, 2005.

[7] S. Kesler, "Mineral supply and demand into the 21st century," *Proceedings for A Workshop on Deposit Modeling, Mineral Resource Assessment, and Their Role in Sustainable Development*, 01 2007.

[8] G. Poxleitner, "Operating costs for miners," SRK Consulting (Canada) Inc., Tech. Rep., 2019. [Online]. Available: https://dxi97tvbmhbca.cloudfront.net/upload/user/image/GPoxleitner_OperatingCostEstimationForMiners_201620191128185443446.pdf

[9] M. Åstrand, M. Johansson, and A. Zanarini, "Underground mine scheduling of mobile machines using constraint programming and large neighborhood search," *Computers Operations Research*, vol. 123, 2020.

[10] M. Schulze and J. Zimmermann, "Staff and machine shift scheduling in a german potash mine," *Journal of Scheduling*, vol. 20, no. 6, pp. 635–656, 2017.

[11] C. Seifi, M. Schulze, and J. Zimmermann, *A two-stage solution approach for a shift scheduling problem with a simultaneous assignment of machines and workers*, 05 2019, pp. 377–385.

[12] M. Nehring, E. Topal, and P. Knights, "Dynamic short term production scheduling and machine allocation in underground mining using mathematical programming," *Mining Technology*, vol. 119, no. 4, pp. 212–220, 2013.

[13] A. Andrade and P. Rampazzo, *Understanding plan's priorities: Short term scheduling optimization*, 05 2019, pp. 386–392.

[14] M. Gamache, P. Cohen, and R. Grimard, "Fleet management system for underground mines," *CIM bulletin.*, vol. 97, no. 1077, p. 66, 2004.

[15] W. Cox *et al.*, "A genetic algorithm for truck dispatching in mining," in *GCAI 2017. 3rd Global Conference on Artificial Intelligence*, ser. EPiC Series in Computing, C. Benzm\"uller, C. Lisetti, and M. Theobald, Eds., vol. 50. EasyChair, 2017, pp. 93–106. [Online]. Available: https://easychair.org/publications/paper/3PFP

[16] N. Vagenas, "Dispatch control of a fleet of remote-controlled/automatic load-haul-dump vehicles in underground mines," *International Journal of Production Research*, vol. 29, no. 11, pp. 2347–2363, 1991.

[17] M. Beaulieu and M. Gamache, "An enumeration algorithm for solving the fleet management problem in underground mines," *Computers Operations Research*, vol. 33, no. 6, pp. 1606–1624, 2006.

[18] W. Cox *et al.*, *A Cooperative Coevolutionary Algorithm for Real-Time Underground Mine Scheduling*, ser. Lecture Notes in Computer Science, 2018, book section Chapter 38, pp. 410–418.

[19] Z. Song *et al.*, "Intelligent scheduling for underground mobile mining equipment," *PLoS One*, vol. 10, no. 6, p. e0131003, 2015. [Online]. Available: https://www.ncbi.nlm.nih.gov/pubmed/26098934

[20] P. Saayman, I. K. Craig, and F. R. Camisani-Calzolari, "Optimization of an autonomous vehicle dispatch system in an underground mine," *Journal of the South African Institute of Mining and Metallurgy.*, vol. 106, no. 2, pp. 77–86, 2006.

[21] A. Moradi Afrapoli and H. Askari-Nasab, "Mining fleet management systems: a review of models and algorithms," *International Journal of Mining, Reclamation and Environment*, vol. 33, no. 1, pp. 42–60, 2017.

[22] C. W. Kim and J. M. A. Tanchoco, “Conflict-free shortest-time bidirectional agv routeing,” *International Journal of Production Research*, vol. 29, no. 12, pp. 2377–2391, 1991.

[23] M. Gamache, M. Beaulieu, and P. Cohen, “Solving the fleet management problem in underground mines by dynamic programming,” in *Proceedings of CAMI 2003, Computer Applications in the Mineral Industries*, S. Fytas, Ed., Chiwetelu, Jan. 2003, pp. –.

[24] H. Anjomshoa *et al.*, “Optimising passing bay locations and vehicle schedules in underground mines,” *Journal of the Operational Research Society*, vol. 64, no. 2, pp. 241–249, 2017.

[25] J. Cohoon, J. Kairo, and J. Lienig, *Evolutionary Algorithms for the Physical Design of VLSI Circuits.* Berlin, Heidelberg: Springer Berlin Heidelberg, 2003, pp. 683–711. [Online]. Available: https://doi.org/10.1007/978-3-642-18965-4_27

[26] J. King *et al.*, “Discrete event control and dynamic scheduling for tele-robotic mining – synopsis,” *CIM bulletin.*, vol. 96, no. 1069, p. 116, 2003.

[27] F. Manríquez, J. Pérez, and N. Morales, “A simulation–optimization framework for short-term underground mine production scheduling,” *Optimization and Engineering*, vol. 21, no. 3, pp. 939–971, 2020.

[28] J. White and J. Olson, “Computer-based dispatching in mines with concurrent operating objectives,” *Min. Eng. (Littleton, Colo.); (United States)*, vol. 38:11, 11 1986.

[29] J. W. White, J. P. Olson, and S. I. Vohnout, “On improving truck/shovel productivity in open pit mines,” *CIM BULLETIN*, vol. 86, no. 973, p. 43, 1993.

[30] C. Bo and H. H. Cheng, “A review of the applications of agent technology in traffic and transportation systems,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 11, no. 2, pp. 485–497, 2010.

[31] R. Erol *et al.*, “A multi-agent based approach to dynamic scheduling of machines and automated guided vehicles in manufacturing systems,” *Applied Soft Computing*, vol. 12, no. 6, pp. 1720–1732, 2012.

[32] S. C. Srivastava *et al.*, “Development of an intelligent agent-based agv controller for a flexible manufacturing system,” *The International Journal of Advanced Manufacturing Technology*, vol. 36, no. 7-8, pp. 780–797, 2007.

[33] D. Weyns and T. Holvoet, "Architectural design of a situated multiagent system for controlling automatic guided vehicles," *International Journal of Agent-Oriented Software Engineering*, vol. 2, no. 1, 2008.

[34] L. Breton, S. Maza, and P. Castagna, "A multi-agent based conflict-free routing approach of bi-directional automated guided vehicles," vol. 2006, 07 2006, p. 6 pp.

[35] L. Branisso *et al.*, "A multi-agent system using fuzzy logic to increase agv fleet performance in warehouses," 11 2013, pp. 137–142.

[36] H. Lin and T. Murata, "Dynamic task assignment of autonomous agv system based on multi agent architecture," *Proceedings of the 2010 IEEE International Conference on Progress in Informatics and Computing, PIC 2010*, vol. 2, 12 2010.

[37] D. Giglio and M. Paolucci, "Agent-based petri net models for agv management in manufacturing systems," vol. 4, 02 2001, pp. 2457 – 2462 vol.4.

[38] G. Icarte, E. Riveros, and O. Herzog, "An agent-based system for truck dispatching in open-pit mines," 01 2020, pp. 73–81.

[39] L. Monostori, J. Váncza, and S. Kumara, "Agent-based systems for manufacturing," *CIRP Annals - Manufacturing Technology*, vol. 55, pp. 697–720, 12 2006.

[40] C. Macal and M. North, *Tutorial on agent-based modeling and simulation*, 01 2014, vol. 4, pp. 11–31.

[41] J.-M. Frayret, "Multi-agent system applications in the forest products industry," *Journal of Science and Technology for Forest Products and Processes*, vol. 1, pp. 15–29, 12 2011.

[42] K. Williamson, *Research Methods for Students, Academics and Professionals: Information Management and Systems*, 2nd ed. Elsevier, 2002.

[43] R. K. Yin, *Case Study Research: Design and Methods.* CA: Beverly Hills: Sage publishing, 1994.

[44] The AnyLogic Company. (2020) Anylogic (version 8.7.3 university). [Online]. Available: https://www.anylogic.com/

[45] O. Ülgen *et al.*, "Simulation methodology: A practitioner's perspective," *Dearborn, MI: University of Michigan*, vol. 1, pp. 7–8, 2006.

[46] R. G. Sargent, "Verification and validation of simulation models," in *Proceedings of the 2010 winter simulation conference.* IEEE, 2010, pp. 166–183.

[47] O. Balci and R. G. Sargent, “A methodology for cost-risk analysis in the statistical validation of simulation models,” *Communications of the ACM*, vol. 24, no. 4, pp. 190–197, 1981.

[48] S. Robinson, *Simulation: The Practice of Model Development and Use.* Chichester, West Sussex, England: John Wiley, 2004.

[49] S. Schlesinger, “Terminology for model credibility,” *Simulation*, vol. 32, no. 3, pp. 103–104, 1979.

[50] M. Rüßmann *et al.*, “Industry 4.0: The future of productivity and growth in manufacturing industries,” *Boston Consulting Group*, vol. 9, no. 1, pp. 54–89, 2015.

[51] D. Weyns and T. Holvoet, “Architectural design of a situated multiagent system for controlling automatic guided vehicles,” *International Journal of Agent-Oriented Software Engineering*, vol. 2, no. 1, pp. 90–128, 2008.

[52] R. G. Smith, “The contract net protocol: High-level communication and control in a distributed problem solver,” *IEEE Computer Architecture Letters*, vol. 29, no. 12, pp. 1104–1113, 1980.

[53] J. D. Thompson, *Organizations in Action.* USA: McGraw-Hill, 1967.

[54] T. W. Malone and K. Crowston, “The interdisciplinary study of coordination,” *ACM Computing Surveys (CSUR)*, vol. 26, no. 1, pp. 87–119, 1994.

[55] T. W. Malone, “Modeling coordination in organizations and markets,” *Management Science*, vol. 33, no. 10, pp. 1317–1332, 1987. [Online]. Available: https://EconPapers.repec.org/RePEc:inm:ormnsc:v:33:y:1987:i:10:p:1317-1332

APPENDIX A AGENT BEHAVIOUR

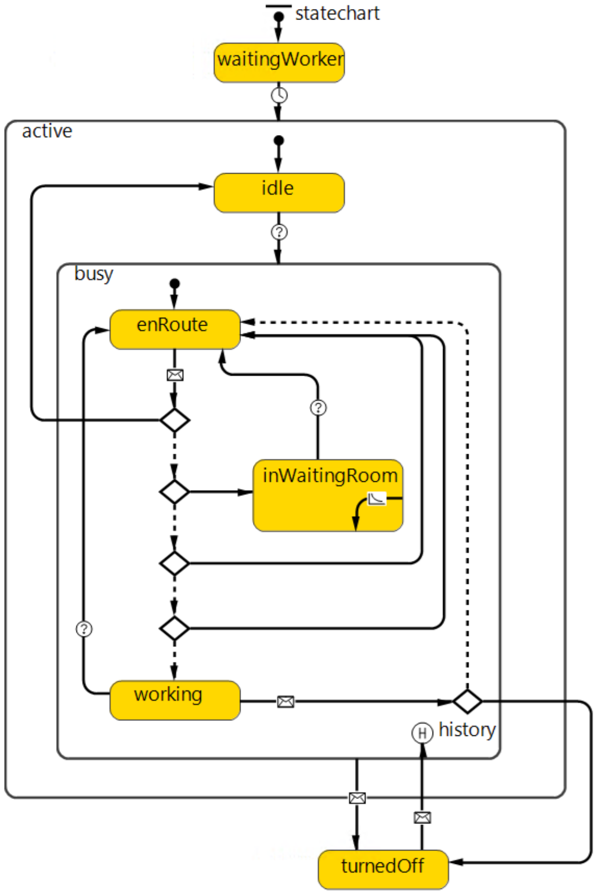

Figure A.1 Statechart of the *equipment* agent (implemented in AnyLogic)

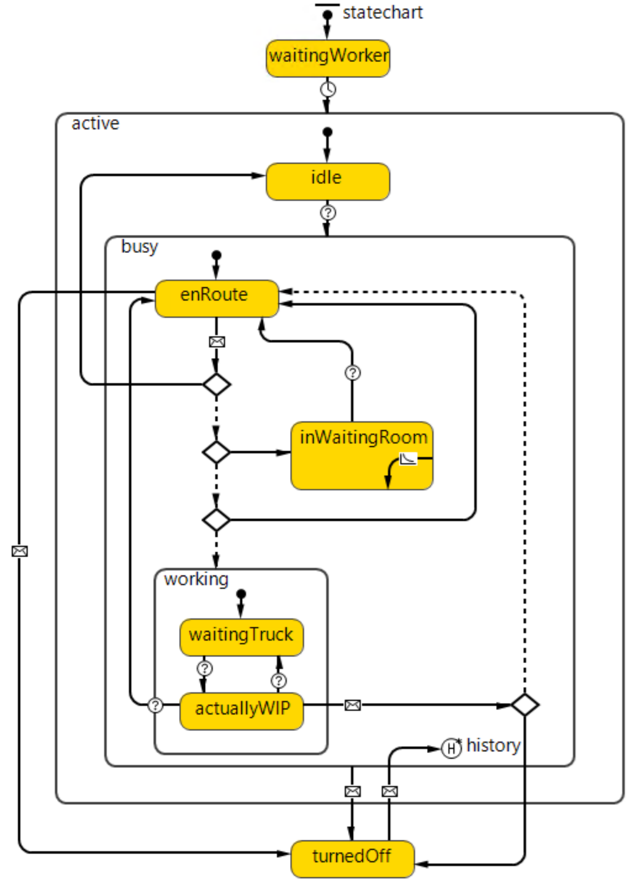

Figure A.2 Statechart of the *LHD* agent (implemented in AnyLogic)

Figure B.1 The user can choose a vehicle from the fleet and climb aboard during the simulation thanks to the on-board camera

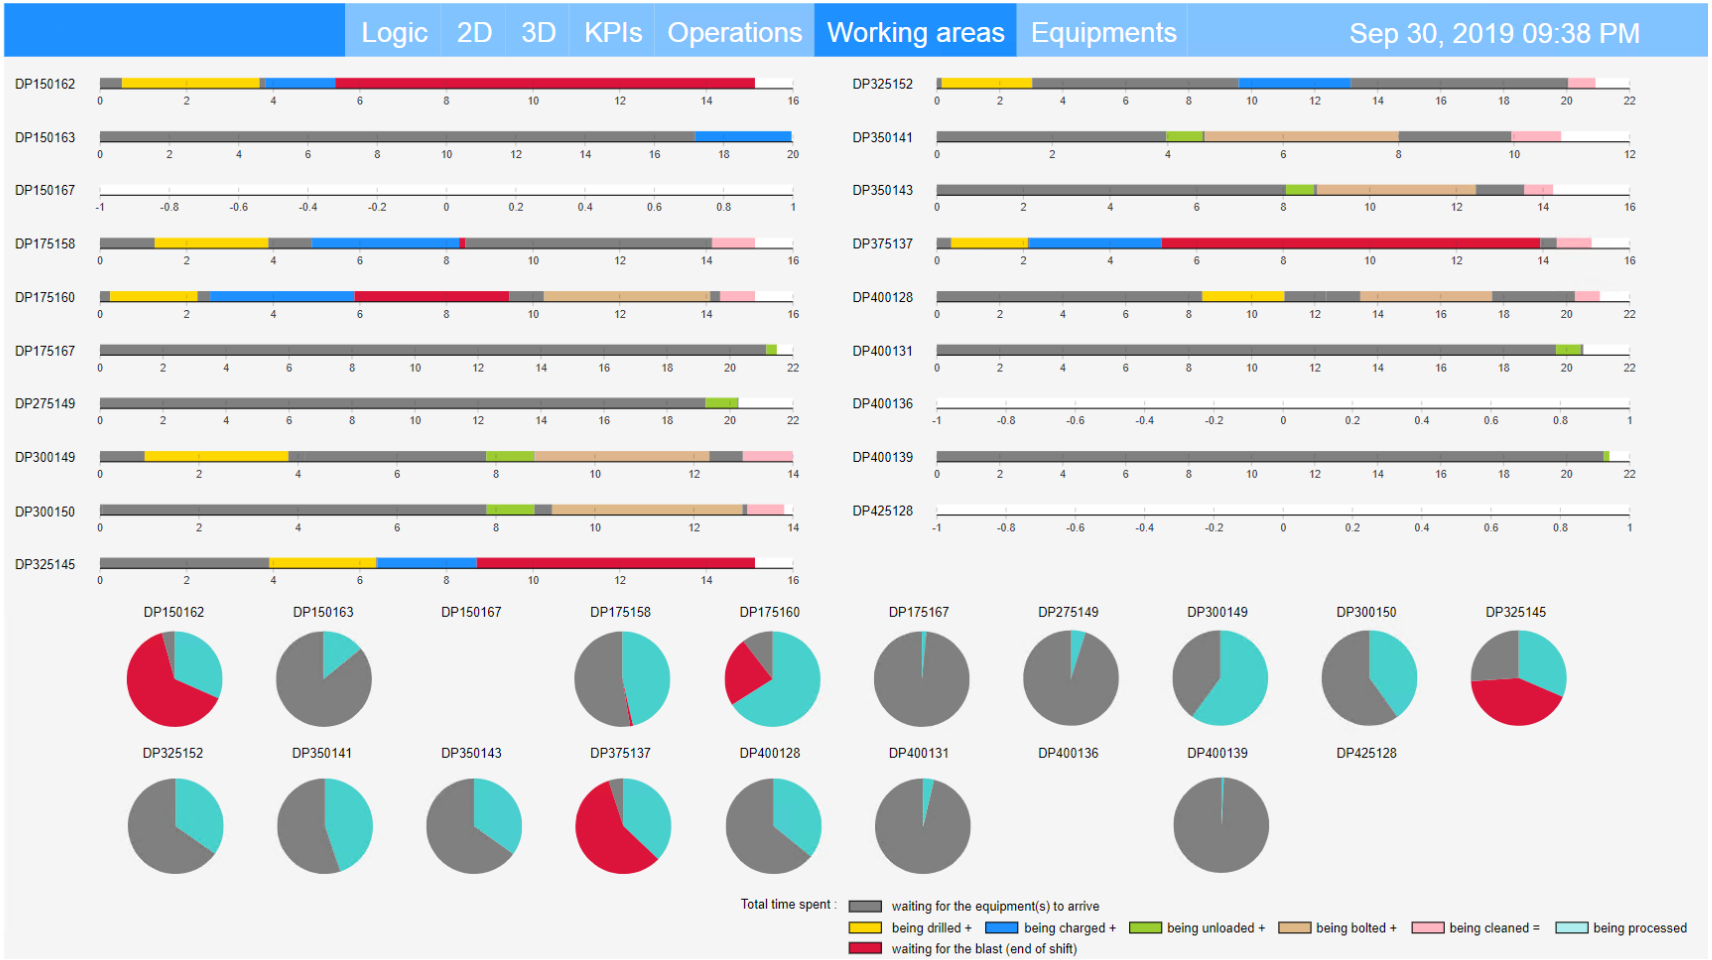

Figure B.2 Example 1 of dashboard accessible during the simulation

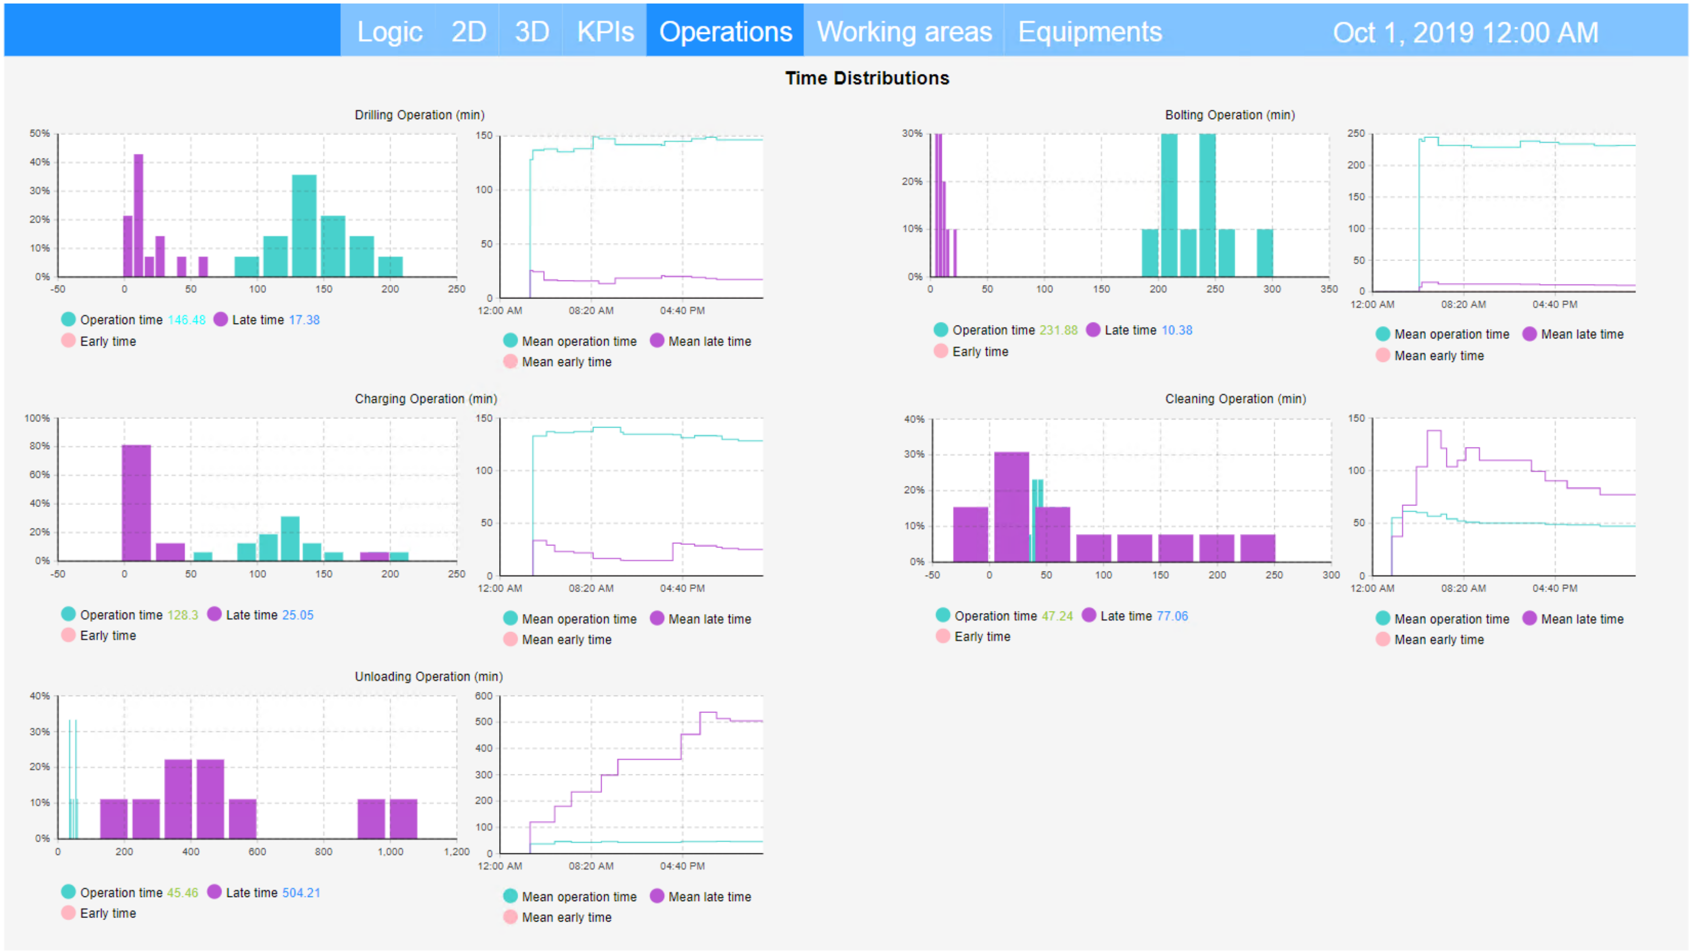

Figure B.3 Example 2 of dashboard accessible during the simulation

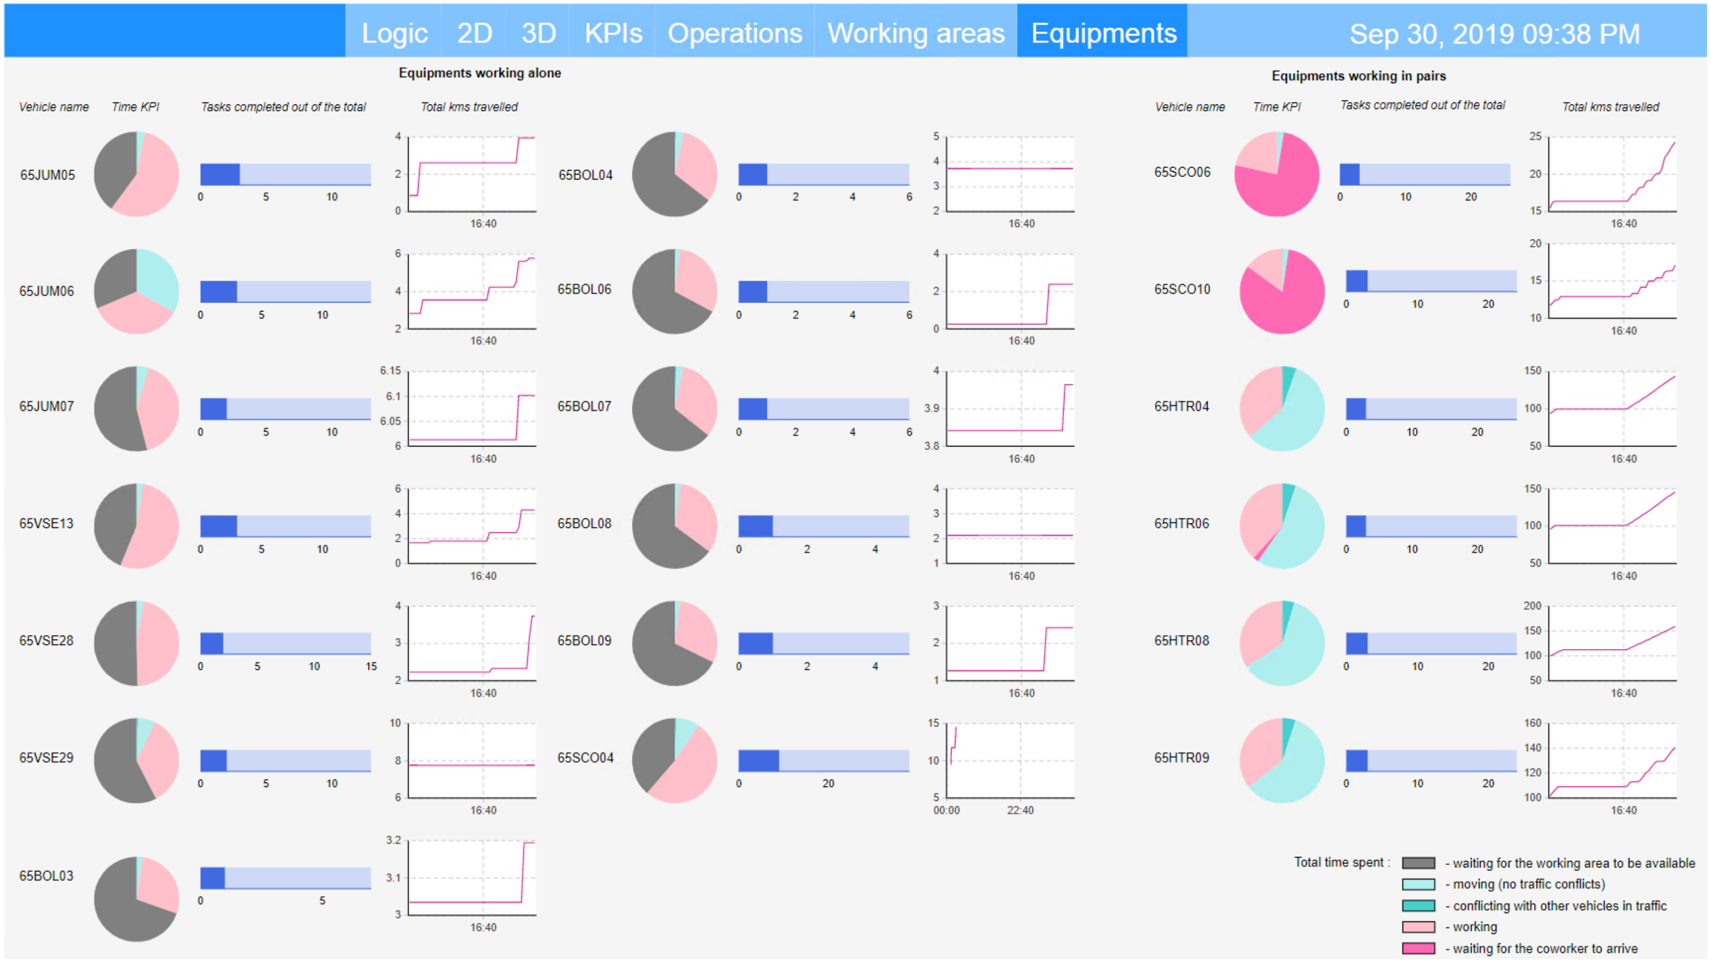

Figure B.4 Example 3 of dashboard accessible during the simulation