Master's Thesis

Automated Planning and Goal Reasoning for Autonomous Systems via Model Checking

Aamir Saeed

A thesis presented for the award of the degree of **M.Sc Mechatronics Engineering**

Submitted to Supervisor: Prof. Michele Taragna



Department of Control & Computer Engineering Politecnico di Torino

July, 2019

Abstract

While AI techniques have found many successful applications in autonomous systems, many of them permit behaviours that are difficult to interpret and may lead to uncertain results. We follow the "verification as planning" paradigm and propose to use model checking techniques to solve planning and goal reasoning problems for autonomous systems. We give a new formulation of Goal Task Network (GTN) that is tailored for our model checking based framework. We then provide a systematic method that models GTNs in the model checker; Process Analysis Toolkit (PAT). We present our planning and goal reasoning system as a framework called Goal Reasoning And Verification for Independent Trusted Autonomous Systems (GRAVITAS) and discuss how it helps provide trustworthy plans in a dynamic environment. Finally, we demonstrate the proposed ideas in an experiment that simulates a survey mission performed by the REMUS-100 autonomous underwater vehicle.

Dedication

I dedicate this dissertation and my project to my loving family and many friends. A special feeling of gratitude to my most loving parents, especially my father (R.I.P) who encouraged me all the time and pushed me for persistence in my work and advised me to be honest always. His words ring in my ears yet. My elder brothers Shakeel, Aqeel and Waqas and my caring sisters Nadia and Sadia were always be with me in any circumstances I could face and have never left my side and are very special.

My dissertation is also dedicated to my many friends who helped me always and supported me throughout my Master's degree. I will always appreciate all they have done, especially Talha for helping me develop my technological skills, Fargham for the many hours of hard work to drag me out if I was stuck in the middle of something during my studies throughout my degree.

I also dedicate my project and give special thanks to one of my good friends and colleagues Hadrien & Zhe for being there for me throughout my entire research work. Both of you have been my best cheerleaders.

Declaration

We, hereby declare that this project neither as a whole nor as a part there of has been copied out from any source. It is further declared that we have developed this project and the accompanied report entirely on the basis of our personal efforts made under the sincere guidance of our supervisors and assistant supervisors. No portion of the work presented in this report has been submitted in the support of any other degree or qualification of this or any other University or Institute of learning, if found we shall stand responsible.

Acknowledgements

In the name of Allah, the Most Gracious and the Most Merciful

Alhamdulillah, all praises to Allah for the strengths and His blessing in completing this thesis. Special appreciation goes to both of my assistant supervisors & mentors, Hadrien & Zhe, for their supervision and constant support. Their invaluable help of constructive comments and suggestions throughout the project and thesis works have contributed to the success of this project. Not forgotten, my appreciation to my supervisor, Jin Song for his support and knowledge regarding this topic.

I would like to express my appreciation to the Consultant, Department of Control & Computer Engineering, and Prof. Michele Taragna (co-supervisor) and also to the Prof. Massimo Violante for their support and help towards my graduate affairs. My acknowledgement also goes to all the lab-technicians and office staffs of Department of Control & Computer Engineering, Politecnico di Torino for their co-operations.

Sincere thanks to all my friends especially Talha, Fargham, Naeem, Asif and others for their kindness and moral support during my graduate studies. Thanks for the friendship and memories.

Last but not least, my deepest gratitude goes to my beloved parents; the late Mr. Saeed-ur-Rahman and Mrs. Hafeeza Begum and also to my brothers and sisters for their endless love, prayers and encouragement. To those who indirectly contributed in this research, your kindness means a lot to me. All of you have been my best cheerleaders.

Contents

1	Intr	oduction	9
	1.1	Planning	10
		1.1.1 Automated Planning	10
	1.2	Goal Reasoning	11
		1.2.1 Goal Seeking	12
		1.2.2 Reasoning	13
	1.3	Model Checking	14
		1.3.1 PAT - An Enhanced Model Checker	15
2	Prel	iminaries with Motivating Example	16
	2.1	Preliminaries	16
	2.2	Motivating Example	19
3	Plan	nning & Goal Reasoning	21
	3.1	Goal Task Networks	21
	3.2	Translating GTN Into CSP#	24
	3.3	Planning Under Resource Constraints	26
	3.4	Goal Reasoning	28
		3.4.1 Reasoning About Rewards/Penalties of Goals	29
		3.4.2 Reasoning About Consistency of Goals	30

	3.5	Performance Testing	32		
4	GRAVITAS				
	4.1	An Overview of GRAVITAS	34		
	4.2	Verification of PAT Models	36		
	4.3	Interacting with Un-trusted Components	37		
5	Imp	lementation & Experiment	39		
	5.1	Implementation of GRAVITAS	39		
		5.1.1 Experimental Setup	39		
	5.2	Simulation in MOOS pMarineViewer	42		
		5.2.1 Primary Stage	42		
		5.2.2 Secondary Stage	45		
6	Rela	ted Work	48		
7	Con	clusion	51		

List of Figures

2.1	The aircraft that disappeared, taking-off back in 2011	19
2.2	An illustration of the overall survey mission	20
4.1	Overall workflow of GRAVITAS	35
5.1	Typical Module Topology	40
5.2	Simulation - Mission started	43
5.3	Simulation shown in MOOS-pMarineViewer.	44
5.4	Simulation - AUV avoids hostile & resumes the mission	45
5.5	Simulation - AUV sleeps & resumes the mission	47

List of Tables

2.1	The syntax and meaning of LTL operators	18
3.1	Performance testing for planning and goal reasoning in two levels of PAT models. Level 1 decides which survey areas to visit and the order to visit them. Level 2 further decides the entry and exit points of each survey area.	33

Chapter 1

Introduction

Artificial intelligence, characterized as intelligence shown by machines, has numerous applications in the modern society. In detailed manner, it is Weak AI, the kind of AI where programs are made to perform explicit tasks, that is being utilized for a wide extent of uses including medical diagnosis, electronic trading and exchanging platforms, robot control, and remote sens-



ing. Artificial intelligence has been utilized to create and propel various fields and enterprises, including finance, health care, training, transportation, and that's only the tip of the iceberg.

This dissertation is inspired by the work of Hadrien Bride, Zhe Hou, Jin Song Dong and few other notable people. In this research, we will apply AI techniques which will allow producing robust ways to deal with modelling behavior, empowering agents' reaction to unforeseen battlefield conditions and improvement of autonomous decision making ability. Mimicked agents in meaningful simulation and virtual conditions ought to be fit for getting objectives and applying intellectual, critical thinking procedures to survey a circumstance and selfsufficiently decide the best way to achieve objectives.

This dissertation will supply sufficient knowledge about Automated Planning and Goal Reasoning which provides a framework for developing methods that will enable autonomous agents capable of problem solving; considering multiple courses of action; coordinating with friendly forces; following chain of command; and using Tactics, Techniques, and Procedures (TTPs) to direct tasks. These agents execute complex undertakings given just mission objectives, beginning/limit conditions, and limitations.

1.1 Planning

Planning is the process of finding a sequence of actions that will achieve a goal. By understanding and anticipating the outcomes of specific actions and acting with intent, the agent can engage in problem solving, which is difficult with scripted behaviors.

Planning is a central and hard computer science problem that is essential in the development of autonomous systems. Many existing solutions require a controlled environment in order to function correctly and reliably. However, there are situations where adaptive autonomous systems are required to run for a long period of time and cope with uncertain events during the deployment. Here we will apply the automated planning techniques that will enable our agent to plan automatically when such uncertain events occur to achieve the specific goal which will lead itself to the completion of it's directed mission.

1.1.1 Automated Planning

Automated planning, sometimes showed as just AI planning, is a kind of artificial intelligence that concerns the affirmation of techniques or activity sequences, typically for execution by intelligent agents, self-governing robots and unmanned vehicles. As opposed to traditional control and order issues, the arrangements are entangled and ought to be found and streamlined in multidimensional space. Automated planning is additionally identified as the decision theory.

Automated planning innovation now plays a vital role in an assortment of requesting applications, ranging from controlling space vehicles and robots to making the game of bridge. These practical applications create new opportunities among theory and practice: seeing what works well in practice leads to better speculations of planning, and better hypotheses lead to better execution of useful applications. Our work is motivated by the requirements of next generation autonomous underwater vehicles (AUV) in law enforcement and defense industries. More particularly, we are currently developing a decision making system suitable for an AUV designed to stay underwater for up to 6 months, with very limited communication with the outside world. The AUV is expected to carry out survey missions on its own and report details of its surveillance at semi-regular intervals. During the mission, the AUV may encounter underwater currents, deep ocean terrain, fishing boats, objects and places of interest, hostile vehicles etc., each of which may affect its ability to achieve its goals. The AUV must be able to decide which goals to pursue when such uncertain events occur and plan tasks to achieve the goals in an agile manner.

In the face of uncertain events in execution, automated planning becomes an even harder problem. In this case, the agent's goal may be affected and thus both selecting a new goal and re-planning are necessary. This generally follows a *note-assess-guide* procedure, where *note* detects discrepancies (e.g., [1]), *assess* hypothesizes causes for discrepancies, and *guide* performs a suitable response [2]. Unlike classical planning where the goal is fixed, when an uncertainty is detected, it is often necessary to change the current goal. This process will be done by the Goal reasoning technique which will be discussed in detail in the following section.

1.2 Goal Reasoning

Goal reasoning concerns the capacity of an intelligent agent to reason about, figure, select, and deal with its goals/destinations. Goal reasoning contrasts from structures in which agents are determined what goals to accomplish, and perhaps how goals can be decayed into sub-goals, yet not how to powerfully and selfsufficiently choose what goals they should seek after. This requirement can be constraining for agents that understand errands in complex conditions when it isn't attainable to physically build/encode total information of what goals ought to be sought after for each possible state. However, in such conditions, states can be reached to an action failure, opportunities can emerge, and occasions can generally occur that unequivocally persuade changing the goals that the agent is currently trying to accomplish.

Furnishing an agent with the capacity to reason about its goals can expand

execution measures for certain applications. Ongoing advances in hardware and software platforms (including the accessibility of fascinating/complex simulators or databases) have progressively allowed the use of smart agents to task that include partially observable and dynamically-updated states (e.g., because of unforeseeable exogenous events), stochastic activities, numerous (coordinating, impartial, or antagonistic) agents, and different complexities. Accordingly, this is a proper time to encourage discourse among researchers with interests in goal reasoning.

Goal reasoning is about selecting a suitable goal for the planning process. There have been various formulations that attempt to solve planning problems in a dynamic environment, including hierarchical planning methods, such as hierarchical task networks (HTN) [3] and hierarchical goal networks (HGN) [4], and goal reasoning systems such as the Metacognitive Integrated Dual-Cycle Architecture (MIDCA) [5] and the goal lifecycle model [6, 7].

In general AI theory, goal reasoning comprises of two parts:

- 1. Goal Seeking
- 2. Reasoning

1.2.1 Goal Seeking

Goal Seeking links the reasoning and planning processes to actual behavior, translating planning operators into actions in the underlying simulation.

Goal seeking is the way toward finding the right input value when just the output is known. The capacity of goal seeking can be incorporated with various types of computer software programs like Microsoft Excel. Goal seeking is a general term used to depict the procedure associated with making sense of your input value dependent on a definitely realized output value. The procedure includes utilizing a particular administrator in a recipe, which can be determined utilizing computer software.

Goal seeking is one of the apparatuses utilized in "what if analysis" on computer software programs. A what if analysis is a process in which investigation is the way toward changing qualities in (Microsoft Excel) cells to perceive how these progressions will influence equation results on the worksheet. When you are goal seeking, you are playing out a what if analysis on a given value, or the output. Thus, fundamentally, you would make a situation by asking "what if the output was X" — or essentially, a cause and effect circumstance.

1.2.2 Reasoning

Reasoning is the act of adding new facts to the agent's knowledge base and using these facts to select the agent's goal. This includes perception of the environment, spatial reasoning, and the application of knowledge rules to determine new facts.

Two types of logical reasoning can be differentiated in addition to formal deduction: induction and abduction. Given a prerequisite or proposition, one can explain the following.

- Deductive reasoning decides if reality of an end can be resolved for that rule, in light of the truth of the premises. Example: "When it rains, things outside get wet. The grass is outside, consequently: when it rains, the grass gets wet." Mathematical rationale and philosophical rationale are ordinarily connected with this sort of reasoning.
- Inductive reasoning endeavors to help an assurance of the rule. It theorizes a rule after various models are taken to be a determination that pursues from a precondition as far as such a rule. Example: "The grass got wet various occasions when it rained, in this manner: the grass dependably gets wet when it rains." While they might be enticing, these contentions are not deductively substantial, see the issue of induction. Science is related with this sort of reasoning.
- Abductive reasoning, a.k.a. surmising to the best clarification, chooses a pertinent arrangement of preconditions. Given a genuine decision and a rule, it endeavors to choose some potential premises that, if true also, can support the conclusion, though not uniquely. Example: "When it rains, the grass gets wet. The grass is wet. In this manner, it may have rained." This sort of reasoning can be utilized to build up a theory, which in turn can be tried by extra reasoning or information. Diagnosticians, detectives, and researchers frequently utilize this kind of reasoning.

Within the context of a mathematical model, the three kinds of reasoning can be described as follows. The construction/creation of the structure of the model is abduction. Assigning values (or probability distributions) to the parameters of the model is induction. Executing/running the model is deduction [8].

1.3 Model Checking

In computer science, model checking or property checking is, for a given model of a framework, comprehensively and naturally checking whether this model meets a given determination. Ordinarily, one has hardware or software systems as a primary concern, though the particular contains security necessities, for example, the nonattendance of gridlocks and comparable basic expresses that can make the framework crash. Model checking is a technique for naturally confirming rightness properties of limited state frameworks.

So as to tackle such an issue algorithmically, both the model of the framework and the determination are defined in some exact scientific language. To this end, the issue is planned as a task in logic, in particular to check whether a given structure fulfills a given legitimate equation. This general idea applies to numerous sorts of logics and appropriate structures. A basic model checking issue is confirming whether a given equation in the propositional logic is fulfilled by a given structure.

Although some of the above formalisms have been successfully applied to solve real life problems, the verification aspect of the problem remains to be addressed. Usually planning is solved by heuristic search, but this approach does not confer a sufficient level of trust. The correctness, safety, and security issues of autonomous systems are particularly important in mission-critical use cases such as our AUV example. To tackle this problem, we turn to formal methods, which have been used to solve planning problems in the literature. For example, Giunchiglia et al. proposed to solve planning problems using model checking [9]; Kress-Gazit et al.'s framework translates high-level tasks defined in linear temporal logic (LTL) [10] to hybrid controllers [11]; Bensalem et al. [12] used verification and validation (V&V) methods to solve planning.

Following the above ideas, we propose a model checking based framework for hierarchical planning and goal reasoning. Model checking is a technology to automatically verify whether certain properties are satisfied by a model using exhaustive search. Model checking is especially strong in addressing uncertain and concurrent behaviours. This technology has been successfully applied to modelling and verifying uncertain environments, such as network attacks which may involve arbitrary behaviour in communication protocols [13].

1.3.1 PAT - An Enhanced Model Checker

PAT (Process Analysis Toolkit) [14] is a self-contained framework for composing, simulating and reasoning of concurrent, real-time systems and other possible domains. It comes with user friendly interfaces, featured model editor and animated simulator. Most importantly, PAT implements various model checking techniques catering for different properties such as deadlock-freeness, divergence-freeness, reachability, LTL properties with fairness assumptions, refinement checking and probabilistic model checking. Besides, we choose PAT because its verification outcome includes a witness trace which can be effectively extracted to form a plan.

Since our planning method is realised in PAT, we can formulate inconsistency and incompatibility as reachability/LTL properties [15] and verify them on the fly. For instance, when a new goal is generated during execution, we can check whether the new goal conflicts with existing goals, and select the subset of goals that are compatible with each other. We can also verify the planning model itself, such that a given planning model does not output plans that may lead to undesired events. Based on the use of PAT, we propose a novel planning and goal reasoning system called Goal Reasoning And Verification for Independent Trusted Autonomous Systems (GRAVITAS), which is amiable to verifiable and explainable plans for autonomous systems. We demonstrate our ideas in a simulation environment which is compatible with modern AUVs.

Chapter 2

Preliminaries with Motivating Example

2.1 Preliminaries

This section is inspired by the related work of Hadrien Bride et al. [16] briefly introduces necessary concepts upon which this work is based. It notably presents the concepts behind model checking as well as PAT – a mature and efficient model checker.

Model checking is an automatic technique for formally verifying finite-state systems. In the context of model checking, specifications of finite-states systems, i.e., properties to be verified by the systems, are often expressed in temporal logic whereas the system to be checked is modelled as a state transition graph. Model checking involves a search procedure which is used to determine whether or not the state-transition graph satisfies the specifications [17].

Models that can be verified using PAT [14] may take several forms, including: CSP# models, timed automata, real-time models and probabilistic models. The latter ones are extensions of the CSP# language.

In this work, we mainly use CSP# – an high-level modelling language supported by PAT. Formally, a CSP# model is a tuple $\langle Var, init_G, P \rangle$ where *Var* is a finite set of global variables, *init_G* is the initial valuation of global variables, and *P* is a process. Variables are typed: either by a pre-defined type (e.g., boolean,

integer, array) or by any user-defined data type. When the type of a variable is not explicitly stated, then, by default, the variable is assumed to be an integer. For instance, an integer variable v and an integer array a can respectively be defined as follows :

var v = 0; $var a[3]: \{0..5\} = [0(3)];$

The range of a variable can be specified. For instance, the annotation ': $\{0..5\}$ ' specifies that the value of each element in *a* must be in the close interval [0,5]. The three values of *a* are initialised with value 0, as denoted by right-hand side of *a*'s declaration – i.e., [0(3)] is equivalent to [0,0,0].

A CSP# process is defined using the following syntax:

 $P(x_1, x_2, ...) = Exp;$

where *P* is the process name, $x_1, x_2, ...$ are the optional parameters of the process, and *Exp* is a process expression, which defines the computation of the process. The running example in this paper uses the following *subset* of CSP#, shown in Backus–Naur form:

1	Exp	::=	Stop Skip	$Ev{Prog} \rightarrow Exp Exp ; Exp Exp Exp$
2			Exp [] Exp	$Exp \iff Exp \mid if (Cond) \{Prog_1\} else \{Prog_2\}$
3			[Cond] Exp	

Let P and Q be CSP# processes. The process expression *Stop* terminate the expression of a process. The *Skip* expression does nothing. The process expression $e\{Prog\} \rightarrow P$ first activates the event labelled by e and executes the statements given by *Prog*, then it proceeds with the execution of *P*. The statements of *Prog* are defined by the syntax and semantics of C# and can therefore manipulate complex data types. The processes P;Q and P||Q respectively express the sequential and parallel composition of processes *P* and *Q*. We use P[]Q to state that either *P* or *Q* may execute, depending on which one performs an event first. On the other hand, P <> Q non-deterministically executes either *P* or *Q*. The expression *if* (*Cond*) *Prog*₁ *else Prog*₂ is self-explanatory. Finally, the expression [*Cond*] *P*, where *Cond* is a boolean expression, defines a guarded process such that *P* only executes when *Cond* is satisfied.

We can check whether a CSP# process *P* satisfies a given specification using the following expression:

 $^{\#} assert P(x_1, x_2, ...) property;$

Formula	mula Meaning	
$\Box f$	f holds in all subsequent states.	
$\Diamond f$	f holds somewhere in the subsequent states.	
$\mathbf{X}f$ f holds in the next state.		
$f_1 U f_2$ f_2 holds somewhere, and before that, f_1 must hold.		
$f_1 \mathbf{R} f_2$	f_2 must hold until f_1 is true; if f_1 is never true, then f_2 holds forever.	

Table 2.1: The syntax and meaning of LTL operators.

where *property* can be *deadlockfree* (the process does progress until reaching a terminating state), *divergencefree* (the process performs internal transitions forever without engaging any useful events), *deterministic* (the process does not involve non-deterministic choices), and *nonterminating* (no terminating states can be reached). Also, we can check whether the transition system can reach a state where a boolean expression *Cond* is satisfied using:

assert P($x_1, x_2, ...$) reaches Cond;

Additionally, we can check whether a process P satisfies a LTL (cf. Huth et al.'s book [18, Section 3.2.1]) formula F using:

```
|| # assert P(x_1, x_2, ...) |= F;
```

The syntax and meaning of LTL operators is given in Table 2.1.

When checking LTL properties, PAT produces a counter-example when the property to be checked cannot be satisfied, and only outputs "yes" when when the property can be satisfied. For reachability properties, which are widely-used in the planning technique of this paper, PAT outputs different information. When the desired states cannot be reached, PAT outputs "no". When the desired states can be reached, PAT produces a witness trace of actions that leads to the desired states. When model checking reachability properties, the user can specify one of the following *verification engines*: If a breath-first search based engine is used and the desired states can be reached, then PAT will output the shortest witness trace, which is useful when finding certain "optimal" plans. Furthermore, the user can tell PAT to output the witness trace that optimises certain criterion. For example, the following code will produce witness traces that respectively yield maximum reward and minimum penalty, assuming that *Cond* is reachable and *reward* and *penalty* are predefined variables:

```
1 #assert P(x_1, x_2,...) reaches Cond with max(reward);
2 #assert P(x_1, x_2,...) reaches Cond with min(penalty);
```

2.2 Motivating Example

One of the common usages of AUVs is surveying underwater areas and reporting back the locations of potential objects of interest. For instance, in the search for MH370, the missing aircraft from Malaysia Airlines Flight 370, AUVs were deployed in deep ocean areas to locate debris of the aircraft [19]. There are also demands and interests from the defence industry to demonstrate the abilities to scan underwater areas for naval mines and dumped arms, as shown in the Wizard of AUS 2018 Autonomous Warrior Trial [20].



Figure 2.1: The aircraft that disappeared, taking-off back in 2011

In this dissertation, we run an example with the following context that demonstrates a common survey mission for the AUV: the AUV is to be deployed at the *initial position* and to be recovered at the *final position*. During the mission, the AUV is expected to scan three *survey areas* and record the locations of *objects of interest* upon identification. Although our technique is general and could be used on all forms of AUVs, we specifically target a torpedo-shaped AUV named REMUS-100 [21], which is equipped with side scanners that are able to detect surrounding objects. The side scanners have a scan range of about 15 meters and therefore, in order to cover large area, the AUV should perform a *lawn mowing* pattern so that the survey area is fully covered. The overall mission is visualized in Figure 2.2.

The challenges of such a survey mission include optimizing the path to minimize energy consumption of the AUV. To deal with uncertainties of the environment, such as changes of survey areas, unexpected events during the transit from



Figure 2.2: An illustration of the overall survey mission.

one location to another, unprecedented appearance of the hostile ships; though our technique must be agile enough to accommodate the dynamics of the environment. The AUV also needs to make smart decisions autonomously, these include the order in which to visit the survey areas and the entry and exit point of each survey area that maximize trajectory efficiency.

Chapter 3

Planning & Goal Reasoning

This chapter discusses how to solve Goal Task Network planning problems using model checking. We build upon the idea of using model checking to solve classical planning problems [22]. Given a classical planning problem containing an initial state, a goal, and a set of actions, one can construct a system model by translating every action into a corresponding state transition in model checking. The initial state of the planning problem is also the initial state of the model by assigning value to each variable accordingly. The goal is expressed using a propositional formula. Then we use model checking to verify the negation of the formula, so that the model checker provides a counterexample path consisting of actions that lead to a state where the formula is satisfied. In this way, a classical planning problem is converted to model checking of the truth of a formula.

Here, we first give a new formalism of the GTN that is suitable for modelling in CSP#. We then propose a model checking based approach to model GTN and solve the planning problem. We also discuss how goal selection – a vital aspect of goal reasoning – can be done in this approach.

3.1 Goal Task Networks

Goal task networks (GTNs) are an extension and unification of hierarchical task networks and hierarchical goal networks [23, 24]. The main conceptual advantage of hierarchical task networks (HTNs), when compared to flat-structured task networks, is their ability to describe dependencies among actions in the form of nested task networks. HTNs have an explicit task hierarchy which generally reflects the hierarchical structure of many real-world planning applications. This hierarchy has decomposition methods which can then be used during the planning phase following the well known *divide and conquer* scheme. Due to this, HTNs planners are much more scalable and performant than classical planners in practice if the hierarchy is well-designed.

Goal task networks of Alford et al. [24] are similar to hierarchical task networks but also consider goals and sub-goals in addition to tasks and sub-tasks. As a result, they inherit the advantages of HTNs but also provide flexibility and reasoning capabilities in goal reasoning. We give an adaptation of the original GTN below with a focus on guarded state transitions, which are in the same form as processes in CSP#.

Let $\mathscr{V} = \{v_0, \dots, v_{d-1}\}$ be a finite set of variables. Without loss of generality, the state *s* of a goal task network over \mathscr{V} is defined as a function $s : \mathscr{V} \to \mathbb{N}$ assigning a non-negative integer to each variable of \mathscr{V} . The set of goal task networks \mathscr{E} is recursively defined as $e \in \mathscr{E} \Leftrightarrow e = \langle E_e, g_e, \tau_e \rangle$ where:

- $E_e \subseteq \mathscr{E}$ is a finite set of sub-tasks/goals,
- $g_e : (\mathscr{V} \to \mathbb{N}) \to \{\bot, \top\}$ is the guard associated with *e*, and
- $\tau_e : (\mathscr{V} \to \mathbb{N}) \to (\mathscr{V} \to \mathbb{N})$ is the state transition function associated with *e*.

Let $e = \langle E_e, g_e, \tau_e \rangle$ be a goal task network. Then *e* can conceptually represent a task or a goal. In the sequel we shall loosely refer to *e* as a task or a goal when the context is clear. When *e* is a task, its guard models the conditions necessary for the task to begin. When *e* is a goal, its guard models the conditions under which the goal is achieved.

Goal task networks whose set of sub-tasks/goals is empty are called *primitive tasks/goals* and describe the elementary block of goal task network executions.

The state of a goal task network evolves during its execution according to the following *firing rules*: A task/goal *e* is *enabled* in state *s* if and only if $g_e(s) = \top$. A task/goal enabled in state *s* can be *fired*, when it does so, it leads to a new state $s' = \tau_e(s)$.

If e is a primitive task/goal then $s \xrightarrow{e} s'$ denote the fact that e is enabled in state s and that its firing leads to state s'. If e is not a primitive task/goal then

 $s \xrightarrow{e} s'$ denote the fact that there exists a *valid execution* of *e* starting in state *s* and leading to state *s'*.

Given an initial state s_0 , a valid execution of e is a sequence e_0, \dots, e_{n-1}, e , of tasks/ goals, where $n \in \mathbb{N}$, such that $\{e_0, \dots, e_{n-1}\} \subseteq E_e$ and $s_0 \xrightarrow{e_0} \dots \xrightarrow{e_{n-1}} s_n \xrightarrow{e} s_{n+1}$. The set of all valid execution starting from a given state s is denoted by Σ_s .

A *GTN planning problem* is tuple $P = \langle e, i \rangle$ where *e* is goal task network and *i* is the initial state of *e*. The set of solutions for *P* is the the set of all valid plans Σ_i , i.e., the set of all valid executions of *e* starting in state *i*.

A formalised model of our GTN definitions in Isabelle/HOL is available online ¹. The following theorem establishes that our GTN formalism can be used to represent the GTN of Alford et al [24]. The other direction is not important here.

Theorem 3.1.1. Given a GTN (I, \prec, α) in Alford et al.'s notation [24] where I is the set of goals and tasks, \prec is a preorder between goals and tasks, and α is a set of labels/names of goal/task instances, there is a corresponding GTN $\langle E, g, \tau \rangle$ in the above definition.

Proof. (Sketch) The above definition simplifies Alford et al.'s notation by not explicitly including names/labels of instances of goals and tasks in the definition. Here we just assume that we need to construct a task/goal in the form of $\langle E, g, \tau \rangle$ for each task/goal in *I* and preserve their relation in the preorder \prec .

Each primitive task/goal in *I* corresponds to a primitive task/goal $\langle E_p, g_p, \tau_p \rangle$ in \mathscr{E} where the E_p part is empty. For each primitive task in *I*, there is a corresponding *operator o* which is a tuple (n, χ_n, e) . Correspondingly, the guard g_p is the precondition χ_n , and the state transition τ_p captures the effect *e* in the operator application.

For a primitive goal release of $\langle E'_p, g'_p, \tau'_p \rangle$, the guard g'_p is the Boolean condition of the goal, and the transition function τ'_p is the identify function.

For each task decomposition using a task method (c, χ_c, gtn_c) , the corresponding task $\langle E_c, g_c, \tau_c \rangle$ has a guard $g_c = \chi_c$, and the E_c part captures the sub-GTN gtn_c . The transition function τ_c is the identity function.

¹https://figshare.com/articles/GTN_thy/6964394

Similarly, for each goal decomposition using a goal method (g, χ_g, gtn_g) , we can construct a goal $\langle E_g, g_g, \tau_g \rangle$ where E_g captures the sub-GTN gtn_g , the guard g_g is the precondition χ_g , and the transition function τ_g is the identity function.

The preorder \prec is captured by set membership relation of tasks/goals and sub-tasks/sub-goals in the *E* part.

3.2 Translating GTN Into CSP#

In a previous work [22], it has been formally established a relation between the classical planning domain and the model checking domain, which helps to reduce errors (e.g., lack of type of information) in manual specification or translation. We provide formal semantics for translating PDDL to CSP# and have implemented the translation in our model checker - Process Analysis Toolkit (PAT). This implementation has been validated with several case studies, which show that using the existing model checker (PAT) to solve classical planning problems is both feasible and efficient.

In this work, we propose a translation from GTNs into CSP# in order to treat goal reasoning within the planning algorithm.

Let $e = \langle E_e, g_e, \tau_e \rangle$ and $\{e_0, \dots, e_{n-1}\} \subseteq E_e$ be GTNs defined over the set of variables $\mathscr{V} = \{v_0, \dots, v_{d-1}\}$. Further, let *i* be the initial state of *e*. The GTN planning problem $P = \langle e, i \rangle$ is modelled as follows.

First, the variables are declared and initialised to their initial values.

```
var v0 = i(v0); \cdots 
var v_{d-1} = i(v_{d-1});
```

Second, the GTN e and its sub-GTNs, as well as their sub-GTNs, are recursively defined using to the following template.

```
 trans_{e}() = [g_{e}] event_{e} \{\tau_{e}\} \rightarrow Skip 
 sub_{e}() = e_{0}() \Leftrightarrow \cdots \Leftrightarrow e_{n-1}() ; (sub_{e}() \Leftrightarrow Skip) 
 se() = sub_{e}() ; trans_{e}()
```

The process $trans_e()$ is guarded by the boolean predicate g_e and, if executed, transforms the current state into its successor state according to the state transition τ_e . Note that τ_e can be effectively encoded by the set of C# statements as C# is Turing complete. Further, the process $sub_e()$ models the set of e's sub-GTNs that

can be executed before executing $trans_e()$. In $sub_e()$, we use non-deterministic choices $\langle \rangle$ to connect the execution of sub-GTNs. The last part of sub_e , i.e., $(sub_e() \langle \rangle skip)$, allows the process to repeat zero or more times, which effectively chooses and executes any sub-GTN zero or more times. This general translation allows the execution of GTN e to be decomposed to the execution of any subset of sub-GTNs $\{e_0, \dots, e_{n-1}\}$ for any number of times. Finally, the CSP# process e() links the processes $sub_e()$ and $trans_e()$ to model the behaviour of the GTN e.

Theorem 3.2.1. For every GTN $\langle E, g, \tau \rangle$, there is a corresponding model in CSP#.

Proof. (Sketch) By the construction of the CSP# process e() and according to the definition of the GTN e, all valid transition sequences of the CSP# process e() correspond to valid plans of the GTN planning problem P.

Example: Take the overall control of the AUV survey mission as an example. At this granularity, the GTN is responsible for making high-level decisions regarding the mission, such as which survey area to visit, in which order, and how to visit it (enter from which direction and exit from which direction). Assuming all the predefined locations are stored in an array, a primitive task at this level is goto(i), which moves the AUV to location i:

```
1 goto(i) = [visited[i] == 0]go.i {

2 currentPosition[0] = position[i][0];

3 currentPosition[1] = position[i][1];

4 visited[i] = 1;

5 } → Skip;
```

In the goto(i) task, the vector *visited*[] records the status of each location. The precondition *visited*[i] == 0 ensures that each location is visited only once. Since goto(i) is a primitive task, it does not contain subtasks/subgoals, therefore, its formulation only involves the guard condition and the transition.

The compound task survey(i) dictates which locations to visit for survey area *i*. This task does not have explicit state transitions, but instead performs state transitions in its subtasks (goto() tasks). Following the translation template, survey(i) is formulated as:

```
survey(i) = (goto(i_0) \Leftrightarrow \cdots \Leftrightarrow goto(i_n)); (survey(i) \Leftrightarrow Skip);
```

where i_0, \dots, i_n are the indices of the locations in survey area *i*.

Similarly, the survey mission is formulated as below:

 $1 \text{ mission}() = (\text{survey}(0) \diamond \cdots \diamond \text{survey}(m)); (\text{mission}() \diamond \text{Skip});$

where $0, \dots, m$ are the indices for survey areas.

The overall GTN involves initialising the start position of the AUV, performing the survey mission, and returning to the final position for recovery. This is modelled as below where we omit the code of *initialise*():

```
1 rendezvous() = goto(finalPosition);
2 main() = (initialise() <> mission() <> rendezvous()); (main() <>
Skip);
```

However, since the motivating example specifies that the three sub-GTNs of *main()* should be executed sequentially, the above definition can be optimised as below:

main() = initialise(); mission(); rendezvous();

3.3 Planning Under Resource Constraints

For most planning applications, considering resource constraints, such as limited amount of available energy, is critical to the quality and relevance of the produced plan. This is particularly true in the application domain we consider as strategic commanders aim at launching AUVs that are meant to operate autonomously for extended period of time with limited resources. Therefore, it is essential that these resource constraints are correctly modelled in order to be able to produce plans that can be fully realised, i.e., plans that do not require more resources than available. Also, as unexpected events may arise during the execution of plans, it is necessary to formulate plans that minimise resource consumption in order to maximise the AUV's resilience.

Suppose we wish to consider a finite set of *m* resources $R = \{r_0, \dots, r_{m-1}\}$ and certain tasks that may consume or produce a finite and discrete amount of one or several of these resources. To do so we introduce, to the GTN modelling of a planning problem, a set of *m* new variables $\mathscr{V}_R = \{v_{r_0}, \dots, v_{r_{m-1}}\}$ modelling the amount of available resources. In the initial state, the values of these variables correspond to the amount of resources available on launch. When a tasks *e* consumes one or several resources, its guard g_e is extended so that it can only be executed if the resources needed to perform it are available before it executes. Additionally,

its state transition function τ_e is also extended so that it decreases the values of the resource variables in order to reflect the resources consumed. Similarly, when a tasks *e* produces one of several resources, its state transition function τ_e is extended so that it increases the values of the resource variables in order to reflect the resources produced.

Example: In the motivating example, we wish to model the AUV energy consumption while moving based on the distance to travel. To do so we introduce the variable *energyLevel* which models the amount of energy left in the battery as well as a function *dist* that returns the distance of the trajectory between two positions and the constant *energyRequiredByMeter* which is used to scale the energy consumption linearly with respect to a travelled distance. We then modify the goto(i) implementation as follows:

```
1 goto(i) = [visited[i] == 0 &&
2 energyLevel >= dist(currentPosition, position[i])] go.i {
3 energyConsumed = dist(currentPosition, position[i]) *
energyRequiredByMeter;
4 energyLevel -= energyConsumed;
5 currentPosition[0] = position[i][0];
6 currentPosition[1] = position[i][1];
7 visited[i] = 1;
8 } → Skip;
```

These changes allow the states of the GTN modelling of a planning problem to encompass available resource quantities and guarantee that valid plans do not, at any time, consume more resources than available. Furthermore, these changes also enable us to minimise resource consumption by maximising the available resource quantities. However, as several resources may be considered, this leads to a multi-objectives optimisation problem that is unfortunately not readily supported by model checking tools like PAT.

We solve this problem by modelling the connections between resources. Note that some resources might be more valuable than others with respect to the mission objectives. Therefore, to avoid the need for multi-objectives optimisation capability, we propose to reduce the problem to a single-objective optimisation. To do so, we suggest the use of an extra variable Λ acting as a *common currency* which is used, among other things, to evaluate the overall state of resources. To update the value of Λ we require, for each resource $r \in R$, a *conversion func*tion $\lambda_r : \mathbb{N} \to \mathbb{N}$ relating the basic unit of a resource as modelled by variable v_r to the basic unit of value of Λ . Conversion functions used in practice include linear functions, logistic functions as well as exponential and logarithmic functions depending on the nature of the resources. Using these conversion functions, we further extend the state transition functions of tasks producing (respectively consuming) resources so that they increase (respectively decrease) the value of Λ accordingly.

An important aspect of this approach is that it enables the comparison of any two sets of quantified resources by transitivity. As a result, maximising the value of Λ minimises the overall resources consumption while accounting for the relative importance of the considered resources. Another important aspect of this approach is that it provides mission operatives with an economic perspective on the complex relations that govern the relative importance of available resources – a familiar perspective people can relate to in everyday life.

To illustrate the use of a conversion function we integrate the common currency into the motivating example by inserting the following line after line 3 of the above code modelling the movement of the AUV, where r_{energy} and e_{energy} are user-defined constant:

```
\Lambda = [r_{energy} * energyConsumed^{e_{energy}}];
```

Continuing with the AUV survey mission example, our model described above already takes the energy cost into account. To find a plan for the modelled GTN with respect to the energy cost, we first need to define the condition for the overall goal:

```
#define goal (∀i. visited[i] == 1) && (currentPosition[0] ==
finalPosition[0] && currentPosition[1] == finalPosition[1]);
```

which states that all the locations are visited, and the AUV's current location is the final position. We then use PAT to find a plan that yields minimal energy cost by model checking the following assertion:

```
\# assert main() reaches goal with max(\Lambda);
```

3.4 Goal Reasoning

In this section we further discuss the concepts that enable our model checking based approach to deal with run-time goal reasoning.

3.4.1 Reasoning About Rewards/Penalties of Goals

Due to environment constraints and resource constraints, the completion of one or several goals may not be possible, or perhaps not worthwhile. Further, goals may not have the same priority. Some goals may be more important to the success of the mission than others. Additionally, as one of the underlying directives is to minimise resources consumption, the produced plans may not consider secondary objectives and only fulfil the minimum requirements in order to complete the mission if the incentive to do so is not correctly modelled.

To cope with these challenges, we propose to associate the achievement of a goal with a reward function relating the goal completion to an amount of the basic unit of value of Λ – the previously introduced variable acting as the common currency. In this setting, maximising the value of Λ prioritises and incentivises the completion of goals providing the most rewards while compromising with the resources they require to be completed. Further, as the resources conversion functions and the reward functions can be arbitrarily complex arithmetic functions, this provides a way to assess trade-offs between complex, competing criteria for a large number of resources and goals.

These economic notions therefore lead to the formulation of highly costeffective plan. Additionally, when multi-agents missions are considered, they provide further benefits as market-based mechanisms [25] can be leveraged to obtain greater collaboration among agents as well as to optimise resources and tasks allocation. These mechanisms also provide non-technical operatives the means to leverage their day to day economic knowledge to specify technical details of the missions that have to be accomplished by the agent.

Example: Returning to the motivating example, we wish to prioritise the recovery of the vehicle (rendezvous()) over the completion of the survey (mission()). To achieve this, we first insert the following code into goto(i) (between the curly braces):

 Λ += reward_{survey};

We then modify the definition of *rendezvous*():

rendezvous() = rend{ Λ += reward_{rendezvous};} \rightarrow goto(finalPosition);

We set *reward*_{rendezvous} to be far greater than *reward*_{survey} × N where N is the total number of positions in the model. We also have to ensure that *reward*_{survey} is greater than $\lfloor r_{energy} * energyConsumed^{e_{energy}} \rfloor$, otherwise PAT will choose not to

visit any position at all. Finally, we modify the *goal* so that visiting all positions and returning to recovery position is no longer mandatory. Rather, we use a more flexible goal, defined as below:

#define goal $\forall i \in C$. visited [i] == 1;

where *C* is a subset of positions that are critical and will override the optimisation on reward/penalty. Now when we model check

```
\# assert main() reaches goal with max(\Lambda);
```

However, if the *energyLevel* is sufficient to visit all positions and go to the recover position, then PAT will output such a plan with minimal energy consumption. Otherwise, if the *energyLevel* is insufficient due to unexpected events such as strong current, energy spent on detour or surveying uncertain objects, etc., PAT will try to find a plan that ensures that the positions in *C* are visited, and that *rendezvous*() is far more likely to be executed than visiting a few more positions.

3.4.2 Reasoning About Consistency of Goals

Consider the following scenario: the AUV has finished the survey mission and now has to report the results. There are two ways to report: (1) acoustic communication with a nearby friendly surface vessel; (2) surface and use satellite communication. Suppose there is no friendly surface vessel nearby, then the AUV will choose the second method. However, suppose there is a hostile surface vessel, which the AUV should avoid. Now the AUV has two goals: report using satellite communication and avoid the hostile surface vessel. The underlying plans for these two goals have conflicts, and the two goals should not be pursued at the same time.

Since PAT can determine whether a condition is satisfiable or not in execution, we can also use PAT to determine the satisfiability of the conjunction of several conditions. To solve the above issue, we first formulate the goals as the conditions below:

```
1 #define goalCompleteSurvey auvCom == 1;
2 #define successfulSurvey goalCompleteSurvey && hvContact == 0;
```

The first goal says that the AUV has done the communication, the second goal is a compound goal that consists of the first goal and that the AUV does not surface when the hostile vessel is nearby (hvContact == 0). We define a task *auvReport*

which consists of subtasks *auvAcousticCom* and *auvSurfaceCom*, which represents communication with friendly surface vessel and with a satellite respectively.

```
auvAcousticCom() = [fvInRange]comFV{auvCom = 1;} → auvReport();
auvSurfaceCom() = [!fvInRange]comS{auvDepth = 0; energyLevel =
10; auvCom = 1; if (hostileInRange) hvContact = 1;} →
auvReport();
auvReport() = auvAcousticCom() [] auvSurfaceCom();
```

The condition *fvInRange* checks whether the friendly vessel is in range for acoustic communication. Verifying the below assertion, which states that *auvReport* can reach a state where the communication has been done and the AUV has not had contact with the hostile vessel, would return negative by PAT.

#assert auvReport() reaches successfulSurvey;

This means that the above two goals are incompatible, and PAT cannot find an execution path to satisfy both. To resolve this issue, we can add a new task that moves the AUV away from the hostile vessel, as coded below:

```
auvAvoidContact() = case {
    hostileInRange: auvMove();auvReport()
    default: auvReport()
};
auvReport() = auvAcousticCom() [] auvSurfaceCom() []
    auvAvoidContact();
```

Now PAT returns affirmative for the above verification and gives a plan to achieve the goal *successfulSurvey*.

Theoretically, we can extend this solution to check incompatibility of a set of goals. Given a set S of goals, we can use PAT as a black-box and implement Algorithm 1 [26] to find the minimal set of goals that are incompatible. We can also find the set of achievable goals, and update the model to resolve unachievable goals if necessary. Algorithm 1 is an elementary method for efficiently finding the minimal unsatisfiable core of a set of formulae by divide and conquer.

Algorithm 1 A simple algorithm for finding minimal unsatisfiable core (MUC). To find a MUC of *S*, call *Minimise*(S, \emptyset).

```
procedure MINIMISE(S, S<sub>0</sub>)

Randomly partition S into two sets S' and S" of the same size.

if S' \wedge S_0 is unsatisfiable then

return Minimise(S', S<sub>0</sub>);

else if S'' \wedge S_0 is unsatisfiable then

return Minimise(S", S<sub>0</sub>);

else > S' \wedge S_0 and S'' \wedge S_0 are both satisfiable

S'_{min} \leftarrow Minimise(S', S_0 \wedge S'');

S''_{min} \leftarrow Minimise(S'', S_0 \wedge S''_{min});

return S'_{min} \wedge S''_{min};

end if

end procedure
```

3.5 Performance Testing

To judge the feasibility and scalability of the model checking based approach, we have tested two levels of planning details. (i) The first level consists of finding an order of the areas to survey so that it minimises the energy cost of the mission. At this level we abstract away the entry, the internal path and the exit point of each survey area. The second level (ii) enables the entry and exit point of each survey area to be determined. These levels respectively correspond to two GTNs of increasing complexity.

We ran the testing on the NVIDIA Jetson TX2 - a power-efficient embedded chip that is equipped in a customised REMUS-100 underwater vehicle at Defence Science and Technology (DST) Australia. We report the results in Table 3.1, in which each configuration is run 5 times and the average of the CPU time and memory usage are displayed. One could theoretically also model the "lawn mowing" path inside each survey area, but it is more of an actuation problem than a planning problem, thus we do not test it here.

The model complexity has a significant impact on the run-time and memory usage of the goal reasoning and planning phase. This is not surprising and is mainly due to the explosion of the state-space size – an issue commonly encountered by model checkers [27]. On the other hand, the REMUS-100 AUV only has a cruising speed of 5.4 km/h, which means that the software has plenty

Level	# of Survey Areas	avg. CPU Time (s)	avg. Memory Usage (MB)
	2	0.005	8.4
1	3	0.01	8.4
1	4	0.03	8.4
	5	0.16	11.7
	2	0.14	11.2
2	3	3.20	66.9
	4	40.26	383.8
	5	290.70	796.3

Table 3.1: Performance testing for planning and goal reasoning in two levels of PAT models. Level 1 decides which survey areas to visit and the order to visit them. Level 2 further decides the entry and exit points of each survey area.

of time to perform re-planning during the mission. The other targeted hardware, the Ocean Glider, is even slower since it relies on water movement to generate forward thrust. We conclude that Level 1 is feasible, and Level 2 is feasible only when the number of survey areas is less than 3. Note that both these levels are high-level operations. We still need to convert high-level operations to low-level operations which can be actuated by the hardware.

The above results highlight the trade-off between performance and guarantees. An approach based solely on model checking is at the moment intractable whereas an approach based solely on heuristics do not provide sufficient guarantees about missions critical elements. Therefore, the above empiric results support the design choice of a hybrid approach for goal reasoning and planning. That is, PAT is suitable for making critical high-level decisions, whereas we need to rely on an external program to translate the high-level plans into low-level plans. The verification of this translation is non-trivial: it includes details such as showing that turning the rudder of the AUV at a certain degree corresponds to going a certain direction in the high-level plan. Such details are hardware-dependent and are not in the scope of this paper. Nonetheless, a carefully designed GTN at an appropriate level of details can, in the context of a hybrid approach, provide better trustworthiness and reliability for the high-level decision-making.

Chapter 4

GRAVITAS

This section describes the *Goal Reasoning And Verification for Independent Trusted Autonomous Systems* (GRAVITAS) – an automated system which enables autonomous agents to operate with trustworthy high-level plans in a dynamic environment.

4.1 An Overview of GRAVITAS

The GRAVITAS framework follows a cyclic pattern composed of four main phases: Monitor, Interpret, Evaluate and Control, which are illustrated in Figure 4.1.

The main operative cycle of GRAVITAS begins with the Monitor (1). This component perceives the environment through the signal processing and fusion of the raw outputs of available sensors. For AUVs, examples of sensors includes accelerometers, gyroscopes, pressure sensors and GPS. It is also in charge of processing these data in order to provide information such as the estimated position and the speed of the agent to the Interpreter (2). This step notably involves techniques such as target tracking which we will not detail here [28]. Once the Interpreter (2) receives the required information, it updates the agent's local model of the system and its environment. This formally defined local model is then forwarded to the Evaluator (3) – a component in charge of assessing the validity of the previously established plan with respect to pre-defined specifications. If the Evaluator assesses the plan as valid, the Controller (5) is tasked with executing the plan. Otherwise, if the Evaluator (3) finds the plan invalid e.g., an uncertain



Figure 4.1: Overall workflow of GRAVITAS.

event creates inconsistencies in the previously established plan and the mission requirements, a new plan needs to be formulated. The formulation of a new plan is accomplished by the joint operation of the Planner and Goals Manager components (4). After a new plan is formulated, the Controller (5) is tasked with executing this plan. This step involves processing based on control theory [29] which we do not discuss here.

The components in the lower loop in Figure 4.1 are orchestrated via the Mission Oriented Operating Suite [30] (MOOS) – a middleware mainly in charge of the communication. The main computational workload of the Evaluator (3), the Planner and the Goal Manager (4) components are powered by PAT. Note that although conceptually the planner and the goal manager are two separated components, in our implementation they are realised in the same PAT model, as discussed

in the examples throughout the previous Chapter. Also, to achieve high efficiency in real-life applications, we use a hybrid approach discussed in Section 3.5.

Also, to achieve high efficiency in real-life applications, we use a hybrid approach to implement planning and goal reasoning: the PAT model performs high-level goal reasoning and planning, and we implement an external actuator to derive a low-level plan from a high-level plan, the former will then be sent to hardware for execution.

4.2 Verification of PAT Models

The key advantage of the model checking based approach is that we can formally verify certain properties for the planning and goal reasoning model. This verification guarantees that the model only permits "correct" high-level plans. Since the verified model is directly used to generate high-level plans in the planning and goal reasoning phase, we can ensure that the generated high-level plans not only are optimised by for max rewards (resp. min penalties), but also are "correct" with respect to the verified properties.

The verification itself is straightforward since the model is already in CSP#. We only need to formulate the properties in the specification language (cf. Chapter 2) and use model checking to verify them.

Example: In the AUV survey example, we are interested in checking whether the model would permit an execution sequence in which the AUV hits an obstacle. The below Boolean condition expresses that the position of AUV does not overlap with any position of obstacles.

```
#define dontRunIntoObstacle (&& index:{0..iNumberOfObstacles 1}@
  (obstacles[index][0] != auvPosition[0] || obstacles[index][1]
  != auvPosition[1]));
```

Using LTL, we can check whether this condition holds *for all* subsequent states in the execution. This is realised by an assertion of the form

 $p \vdash \Box c$

where p is a process in CSP#, c is the condition we need to check, and \Box is a modality in LTL that means c holds for all subsequent states. The above verification is realised in the code below:

```
# assert main() |= [] dontRunIntoObstacle;
```

and PAT can automatically return "yes" as the result. Thus we obtain the following lemma:

Lemma 4.2.1. *The example planning and goal reasoning model described in the previous Chapter does not generate plans where the AUV runs into any obstacle.*

Using the same technique, we have verified the following lemmas:

Lemma 4.2.2. The example model described in the previous Chapter does not generate plans where the AUV runs out of battery during the mission.

Lemma 4.2.3. The example model described in the previous Chapter does not generate plans where the AUV surfaces at a location within 3 units of distance of a hostile vessel.

4.3 Interacting with Un-trusted Components

Although the Level 2 planning and goal reasoning model in Table 3.1 suffices in our demonstration of the AUV survey mission, there might be other applications where model checking cannot provide detailed plans in time. For instance, the user may need to adopt heuristic-based planning techniques for UAVs and land vehicles because they run faster.

Inspired by Clarke et al.'s counterexample-guided abstraction refinement [31], we propose to integrate heuristic-based planning techniques as an "un-trusted component" as follows: We treat the heuristic method as an high-level plan generator. Whenever the heuristic method generates a plan, we simulate this plan using the corresponding high-level planning and goal reasoning model, i.e., the *CSP# model*, in PAT. This simulation is much faster than model checking because we only need to check one path of actions instead of checking all paths. If the simulation is successful, then this plan is in the set of plans that can be generated by the CSP# model. If the CSP# model has been verified as described in Section 4.2, then this plan is correct with respect to the verified properties. If the simulation fails, then we add the old plan into a set of *disabled plans* and constraints the heuristic method such that it does not generate one of the disabled plans. This procedure provides plans that have the same formal guarantee as those generated by PAT, but

this procedure may not yield optimal plans. Nonetheless, this procedure provides the means to interact with existing heuristic-based planning techniques generally employed without safe-guards in a reliable way.

Chapter 5

Implementation & Experiment

5.1 Implementation of GRAVITAS

In situ, experimentation is very expensive and slow. While it is mandatory to the final evaluation of the implementation, in this paper we focus on assessing and demonstrating the feasibility of the proposed goal reasoning and planning approach in a virtual environment. Notable challenges include controlling the complexity of the GTN so that the embedded hardware of the AUV is able to carry the computational load in a reasonable time (i.e., less than a minute) and the transposition of a discrete plan as issued by PAT into its continuous counterpart so that it can be enacted by the AUV.

We have implemented the proposed approach and integrated it within a virtual environment closely simulating the mission described in Section 2.2. We first introduce the integration of PAT within a *community* of MOOS applications [30]. We then report the obtained results and discuss the conclusion drawn from them.

5.1.1 Experimental Setup

The experimental setup as shown in the Figure 5.1 is composed of three MOOS communities - *shoreside community*, which simulates the computer at the shore side to communicate with the AUV through sattelite communication, *AUV community*, that simulates the AUV behaviors in the virtual environment and *Hostile community*, which simulates the Hostile which is roving around throughout the



Figure 5.1: Typical Module Topology

mission. These communities include the following:

- *MOOSDB*: a central server application for all communications.
- *pMarineViewer:* a GUI rendering vehicles and associated information during operation or simulation.
- *pShare:* allows data to be shared between MOOS communities.
- *pHostInfo:* determines the IP address of the machine on which it is running and post it to the MOOSDB.
- *uFldShoreBroker:* a tool for brokering connections between a shoreside community and one or more nodes (simulated or real vehicles).
- *uFldNodeBroker:* a tool for brokering connections between a node (a simulated or real vehicle) and a shoreside community.
- *uSimMarine:* a 3D vehicle simulator that updates vehicle state, position and trajectory, based on the present actuator values and prior vehicle state.

- *pLogger:* records the activities of a MOOS session.
- *pMarinePID:* a PID controller for heading, speed and depth.
- *pHelmIvP*: publishes a steady stream of information that drives the platform, typically regarding the desired heading, speed or depth.
- *pBasicContactMgr:* posts to the MOOSDB summary reports about known contacts (detected by sensors) after getting information from the sensors.
- *pProcessWatch:* monitors the health of a set of MOOS application.
- uFldMessageHandler: handles incoming messages from a remote MOOSDB.
- *uXMS*: a terminal based MOOS app for scoping the MOOSDB.
- *pHostileDetect:* integrates PAT and provide goal reasoning and planning ability to the AUV and detects Hostile.

Key MOOS application & it's role

The MOOS application *pHostileDetect* implements the GRAVITAS framework as described in Chapter 4. It subscribes to and monitors channels which broadcast information about the general state of the AUV (e.g., position, speed, heading) as well as information about the objects detected by the side scanners. Then, at each iteration of its internal loop, it interprets this information and models a local world view of the environment. Based on this internal representation and according to the proposed planning approach, it evaluates the actual plan being enacted and, if required, updates it before enacting it by publishing the desired heading, speed and depth of the AUV to the community.

The plan issued by PAT as a part of the re-planning step is a discrete sequence of primitive tasks (e.g. go to 3D position) that require some processing in order to be enacted by the AUV as actuators commands (e.g., set heading, set speed). For instance, the trajectory between several way-points set by the plan has to be compliant with the maximum turn-rate of the AUV. To solve this issue, as a proof of concept we implemented an algorithm based on piecewise Bezier curves composition with continuous curvature constraint for continuous path planning [32]. In the future we plan on using a more advanced low-level planning approach such as the FMT* algorithm [33] that will enable us to consider trajectories based on 3D current dynamics as well as uncertainty in the AUV position.

Primarily, AUV was only capable of detecting the strong currents and mines during the mission, but after the extension of AUV's perception to the environment, it is also capable of detecting hostile ships and performs respective actions in the diverse environment and resumes the mission after being undetected.

Finally, we can demonstrate the proposed ideas in an experiment that simulates a survey mission performed by the REMUS-100 autonomous underwater vehicle.

5.2 Simulation in MOOS pMarineViewer

We demonstrate a case study scenario in a simulation in MOOS. In this scenario, we intend to capture GRAVITAS's capabilities in dealing with dynamic events during execution. We create a survey mission similar to Figure 2.2. Note that although the following screenshots are in 2D, the simulation is actually in a 3D environment.

We set 3 survey areas: lower-left (LL), upper-right (UR), and lower-right (LR), with rewards 22807, 51918, 31313 respectively. Initially, the AUV has an energy level of 60000. During execution, we randomly generate a strong water current, with a chance of 20%, that doubles the energy consumption for an uncertain period of time. For simplicity in this example, we trigger goal reasoning and re-planning at the end of each survey area, although these can be done as frequently as required provided that the computation does not take longer than the interval between re-plannings.

5.2.1 Primary Stage

Figure 5.3a shows the initial plan computed by GRAVITAS. The numbers indicate the high-level plan computed by PAT and the dots indicate the low-level plan generated by the actuator. That is, PAT finds the optimal order as well as the entry and exit points for the survey areas, and the actuator computes a smooth path that



Figure 5.2: Simulation - Mission started

the AUV can follow. The "wing" of the AUV indicates the coverage of the side scan sonar.

In this run, the random generator creates a strong current during the first survey. As a result, the expected energy consumption for the first survey is 14400, but the actual consumption is 23284. Consequently, the Interpreter in GRAVITAS uses a simple "learning" to update the expected energy consumption for future execution to 162% of the estimation. However, this unexpected change causes re-planning in which PAT decides that there is insufficient energy to complete 3 survey areas, and then finds a new plan to optimise the outcome, as shown in Figure 5.3b. In the new plan, PAT chooses to only survey area LR because it yields more reward.

During the transit to the second survey, the water current has returned to normal. At the end of the second survey, the Controller in GRAVITAS discovers that survey area LR has been fully covered before going through the last pass in the





(c) 2nd re-planning.

(d) Mission finished successfully.

Figure 5.3: Simulation shown in MOOS-pMarineViewer.

"lawn-mowing" pattern. Therefore it triggers re-planning and PAT and the actuator enact a new plan, shown in Figure 5.3c, to directly go to the rendezvous point. At the same time, the Interpreter captures that the expected energy consumption for the second survey is 27216, but the actual consumption is 17411, so it lowers the scale of future energy consumption to 104% of the estimation. There is no more strong current on the way to the rendezvous point, and the expected energy consumption is roughly the same as the actual value, and the AUV successfully finishes the mission, as shown in Figure 5.3d. In this case study, the (re-)planning takes around 1 second, which is fast enough for the operation of AUV.

5.2.2 Secondary Stage



(c) AUV keeps the track

(d) AUV resumes mission

Figure 5.4: Simulation - AUV avoids hostile & resumes the mission

If AUV detects any hostile ship during the mission, our environment enables two kinds of actions:

- Avoid Hostile: avoids the hostile suddenly after detecting it and change the heading until the hostile is not in it's mission territory and resumes the mission afterwards as shown in Figure 5.4.
- *Wait:* turns off the engine after updating the speed to 0 & enables the stealth mode if the hostile is close enough and there is a fear of being detected and then resumes the mission after confirming that the hostile is not in the AUV's territory anymore as shown in Figure 5.5.



(c) AUV is in stealth mode

(d) AUV resumes the mission

Figure 5.5: Simulation - AUV sleeps & resumes the mission

Chapter 6

Related Work

Different approaches exist according to the assumptions about the domain, the goals, the plans and the planning algorithm. Conceptually, the domain evolves according to the performed actions, a controller provides the actions according to the observations on the domain and a plan [34]. An example of applying automated reasoning techniques on planning is Kress-Gazit et al.'s framework which automatically translates high-level tasks defined in linear temporal logic formulae to hybrid controllers [11]. This framework allows for reactive tasks, which may change depending on the information the robot gathers at runtime. This is similar to the goal reasoning literature where goals may change depending on the environment at runtime.

This work follows the "planning as model checking" paradigm, which dates back to 1990s, e.g., in the work by Giunchiglia and Traverso [9]. They proposed to solve (classical) planning problems model-theoretically, where planning domains are formalised as semantic models, properties of planning domains are formalised as temporal formulae, and planning is done by verifying whether temporal formulae are true in a semantic model. This idea has been studied and improved in their subsequent work [35, 36, 37], which involves using Binary Decision Diagram based heuristic symbolic search. Similar ideas have been used in planners such as MIPS [38], which can effectively handle the STRIPS subset of the PDDL, and some additional features in ADL.

Closely related to the above work is the verification and validation (V&V) based method of Bensalem et al. [12]. They argue that constructing correct and reliable planning systems is error-prone due to the non-deterministic nature of

planning problems, thus it is important to develop V&V methods for planners to ensure that the generated plans are correct. To achieve this, the authors proposed to *use* V&V techniques to perform planning, and use planning to perform V&V. This work is similar in the sense that we are using model checking techniques to perform planning, and since the planning system is built upon the model checker, we can also verify correctness and safety issues of the plans and goals. As a result, we can not only output plans that are efficient in certain criteria, but also those that are verified safe and correct, which is essential in building trusted intelligent agent and is often required in mission-critical operations.

Goal reasoning has been used in a number of projects about controlling autonomous machines in a dynamic environment. Many goal reasoning systems follow a *note-assess-guide* procedure, and extend it with a cycle of executions to handle the dynamics of the environment and perform goal reasoning and re-planing on-the-fly. Cox et al. [5] propose to use classical planning to formalise goal reasoning. They present an architecture with a cognitive layer and a metacognitive layer to model problem-solving and dynamic event management in self-regulated autonomy. The architecture is realised in the Metacognitive Integrated Dual-Cycle Architecture (MIDCA) version 1.3, which is shown useful in experiment. A detailed account is given by Dannenhauer [39].

Roberts et al. [6] give more detailed definitions of goal reasoning in their framework. They divide the states and goals into two parts: the external part is a modified or incomplete version of the transition system, and the internal part represents the predicates and state required for the refinement strategies. The authors use a data structure called *goal memory* to represent the relationship between goals, subgoals, parent goals etc., and propose to solve the goal reasoning problem using refinement. They use a *goal lifecycle* model to capture the evolution of goals and the decision points involved in the process. The goal lifecycle includes the formulation, selection, expansion, execution, dispatch, evaluation, termination, and discard of goals. This model is adapted by Johnson et al. [7], who give a system called Goal Reasoning with Information Measures. In the scenario of controlling Unmanned Air Vehicles to survey certain areas, the goals are formulated with parameters such as *maximum uncertainty* in the search area, *acceptable uncertainty* under which the goal is considered complete, and *deadline* by which the search must complete. The goal reasoning method is shown useful for unmanned aerial vehicles operating in dynamic environments.

A more theoretical foundation about planning and goal reasoning is sur-

veyed by Alford et al. [24]. The authors unify HGN planning and HTN planning into GTN planning. They also provide plan-preserving translations from GTN problems to HTN semantics. Several computability and tractability results are given. For example, GTN, HTN, and HGN are semi-decidable, and a restricted form called GTN_{*I*} is NEXPTIME. An application of HTN planning realised by symbolic model checking is presented by Kuter et al. [40]. While their work is focused on the theoretical foundation of the problem and they assume fullobservability, this paper is more concerned with a more concrete real-life problem: the execution of the AUV in an uncertain environment. Thus this paper is more focused on practical issues that arise when solving the AUV survey problem.

One interesting use case of goal reasoning is goal selection. Rabideau et al. [41] give a tractable goal selection method algorithm specialised for selecting goals at runtime for re-planning in a system where computational resources are limited and the complete goal set oversubscribe available resources. Kondrakunta and Cox [42] also consider the situation where an agent has more goals than can complete in a given time constraint and show how an intelligent agent can estimate the trade-off between performance gains and resource costs. Another important aspect of goal reasoning is to detect inconsistency or incompatibility of goals and plans. Tinnemeier et al. [43] propose a mechanism to process incompatible goals which have conflicting plans. They argue that the agent should not pursue goals with conflicting plans, and their mechanism can help the agent choose from incompatible goals.

An important application of our project is applying the planning and goal reasoning framework to AUVs. Among many relevant papers, goal reasoning for AUVs [44] is particularly interesting. The authors use a goal-driven autonomy conceptual model which has three parts: the planner, the goal controller, and the state transition system. The goal reasoning problem is formalised in PDDL, which is the standard language for representing classical planning problems and is widely used by many planners. The authors test their approach in simulations where the AUV surveys a defined area and it has to respond (change the goal) to the actions from a nearby unmanned surface vehicle dynamically. Cashmore et al.'s work [45] describes a planning algorithm for AUVs. Like many other related papers, their work assumes certain requirements that are slightly different from our settings. For example, they are focused on temporal planning with time constraints whereas our mission does not have such constraints.

Chapter 7

Conclusion

This work describes a decision-making framework named GRAVITAS for autonomous systems. The GRAVITAS framework not only provides theoretical foundation for hierarchical planning and goal reasoning, i.e., modelling GTNs using CSP# and using model checking to perform planning and goal reasoning, but also includes practical implementations via the model checker PAT and the MOOS application community. This framework is ultimately realised on a hardware chip that runs on the REMUS-100 AUV. Our simulation has shown that the model checker PAT is sufficient to perform high-level decision making tasks. We have also developed various auxiliary functionalities in GRAVITAS to extend the high-level PAT plan into low-level plans for actuation. An important future work is to improve the level of trustworthiness by extending the verification from high-level plans to low-level plans. We are also planning to conduct more realistic simulations, and will attempt to show in situ ability of our approach in real-world demonstrations.

Bibliography

- Matthew Paisner, Michael T. Cox, and Don Perlis. Symbolic anomaly detection and assessment using growing neural gas. In *Proc. 25th IEEE International Conference on Tools with Artificial Intelligence*, pages 175–181. IEEE Computer Society, 2013.
- [2] Michael L. Anderson, Tim Oates, Waiyian Chong, and Donald Perlis. The metacognitive loop I: enhancing reinforcement learning with metacognitive monitoring and control for improved perturbation tolerance. *J. Exp. Theor. Artif. Intell.*, 18(3):387–411, 2006.
- [3] Kutluhan Erol, James A Hendler, and Dana S Nau. UMCP: A sound and complete procedure for hierarchical task-network planning. In *AIPS*, volume 94, pages 249–254, 1994.
- [4] Vikas Shivashankar, Ugur Kuter, Dana Nau, and Ron Alford. A hierarchical goal-based formalism and algorithm for single-agent planning. In Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2, pages 981–988. International Foundation for Autonomous Agents and Multiagent Systems, 2012.
- [5] Michael T. Cox, Zohreh Alavi, Dustin Dannenhauer, Vahid Eyorokon, Hector Munoz-Avila, and Don Perlis. MIDCA: A metacognitive, integrated dual-cycle architecture for self-regulated autonomy. In AAAI, pages 3712– 3718, 2016.
- [6] Mark Roberts, Thomas Apker, Benjamin Johnston, Bryan Auslander, Briana Wellman, and David W. Aha. Coordinating robot teams for disaster relief. In Proceedings of the Twenty-Eighth International Florida Artificial Intelligence Research Society Conference, FLAIRS 2015, Hollywood, Florida. May 18-20, 2015., pages 366–371, 2015.

- [7] Benjamin Johnson, Mark Roberts, Thomas Apker, and David W. Aha. Goal reasoning with informative expectations. In *ICAPS Workshop*, London, UK, 2016.
- [8] Jo Reichertz. Induction, deduction, abduction. In Uwe Flick, editor, *Qualitative Data Analysis*. SAGE Publishing, 2014.
- [9] Fausto Giunchiglia and Paolo Traverso. Planning as model checking. In *European Conference on Planning*, pages 1–20. Springer, 1999.
- [10] Amir Pnueli. The temporal logic of programs. In *Proceedings of the 18th Annual Symposium on Foundations of Computer Science*, SFCS '77, pages 46–57, Washington, DC, USA, 1977. IEEE Computer Society.
- [11] H. Kress-Gazit, G. E. Fainekos, and G. J. Pappas. Temporal-logic-based reactive mission and motion planning. *IEEE Transactions on Robotics*, 25(6):1370–1381, 2009.
- [12] Saddek Bensalem, Klaus Havelund, and Andrea Orlandini. Verification and validation meet planning and scheduling, 2014.
- [13] Guangdong Bai, Jike Lei, Guozhu Meng, Sai Sathyanarayan Venkatraman, Prateek Saxena, Jun Sun, Yang Liu, and Jin Song Dong. Authscan: Automatic extraction of web authentication protocols from implementations. In NDSS, 2013.
- [14] Jun Sun, Yang Liu, Jin Song Dong, and Jun Pang. PAT: Towards flexible verification under fairness. volume 5643 of *Lecture Notes in Computer Science*, pages 709–714. Springer, 2009.
- [15] Edmund M. Clarke, E Allen Emerson, and A Prasad Sistla. Automatic verification of finite-state concurrent systems using temporal logic specifications. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 8(2):244–263, 1986.
- [16] Hadrien Bride and Zhe Hou and Jin Song Dong et al. Planning and Goal Reasoning for Autonomous Systems via Model Checking.
- [17] Edmund M. Clarke, Jr., Orna Grumberg, and Doron A. Peled. *Model Check-ing*. MIT Press, Cambridge, MA, USA, 1999.

- [18] Michael Huth and Mark Ryan. Logic in Computer Science: Modelling and Reasoning About Systems. Cambridge University Press, New York, NY, USA, 2004.
- [19] news.com.au. New ocean infinity search for mh370 encountering big problems. http://www.news.com.au/travel/travel-updates/incidents/newocean-infinity-search-for-mh370-encountering-big-problems/newsstory/89efc71d393585bb3efc1fcd10b765aa, 2018.
- [20] Defence Science Institute. Wizard of aus 2018 an autonomous warrior 2018 trial. http://www.defencescienceinstitute.com/2017/11/06/wizard-aus-2018autonomous-warrior-2018-trial/, 2018.
- [21] Woods Hole Oceanographic Institution. Remus 100. http://www.whoi.edu/main/remus100, 2018.
- [22] Yi Li, Jin Song Dong, Jing Sun, Yang Liu, and Jun Sun. Model checking approach to automated planning. *Formal Methods in System Design*, 44(2):176–202, 2014.
- [23] Vikas Shivishankar. Hierarchical Goal Network Planning: Formalisms and Algorithms for Planning and Acting. PhD thesis, PhD thesis, Department of Computer Science, University of Maryland College Park, 2015.
- [24] Ron Alford, Vikas Shivashankar, Mark Roberts, Jeremy Frank, and David W. Aha. Hierarchical planning: Relating task and goal decomposition with task sharing. In *IJCAI*, pages 3022–3029, 2016.
- [25] Scott H Clearwater. *Market-based control: A paradigm for distributed resource allocation.* World Scientific, 1996.
- [26] Isil Dillig. Lecture notes on automated logical reasoning. https://www.slideshare.net/pvcpvc9/lecture17-31382688, Accessed 2018.
- [27] Antti Valmari. The state explosion problem. In Advanced Course on Petri Nets, pages 429–528. Springer, 1996.
- [28] Subhash Challa, Mark R. Morelande, Darko Mušicki, and Robin J. Evans. *Fundamentals of Object Tracking*. Cambridge University Press, 2011.

- [29] Ernest Bruce Lee and Lawrence Markus. Foundations of optimal control theory. Technical report, Minnesota University Minneapolis Center for Control Sciences, 1967.
- [30] Paul Michael Newman. MOOS-mission orientated operating suite. 2008.
- [31] Edmund Clarke, Orna Grumberg, Somesh Jha, Yuan Lu, and Helmut Veith. Counterexample-guided abstraction refinement. In E. Allen Emerson and Aravinda Prasad Sistla, editors, *Computer Aided Verification*, pages 154– 169, Berlin, Heidelberg, 2000. Springer Berlin Heidelberg.
- [32] Ji-Wung Choi, Renwick Curry, and Gabriel Elkaim. Piecewise bezier curves path planning with continuous curvature constraint for autonomous driving. In *Machine learning and systems engineering*, pages 31–45. Springer, 2010.
- [33] Lucas Janson, Edward Schmerling, Ashley Clark, and Marco Pavone. Fast marching tree: A fast marching sampling-based method for optimal motion planning in many dimensions. *The International journal of robotics research*, 34(7):883–921, 2015.
- [34] Malik Ghallab, Dana Nau, and Paolo Traverso. *Automated Planning and Acting*. Cambridge University Press, New York, NY, USA, 1st edition, 2016.
- [35] Alessandro Cimatti and Marco Roveri. Conformant planning via symbolic model checking. J. Artif. Intell. Res. (JAIR), 13:305–338, 2000.
- [36] Alessandro Cimatti, Enrico Giunchiglia, Fausto Giunchiglia, and Paolo Traverso. Planning via model checking: A decision procedure for ar. In *European Conference on Planning*, pages 130–142, 1997.
- [37] Piergiorgio Bertoli, Alessandro Cimatti, and Marco Roveri. Heuristic search+ symbolic model checking= efficient conformant planning. In *IJCAI*, pages 467–472. Citeseer, 2001.
- [38] Stefan Edelkamp and Malte Helmert. Mips: The model-checking integrated planning system. *AI magazine*, 22(3):67, 2001.
- [39] Dustin Dannenhauer. *Self monitoring goal driven autonomy agents*. PhD thesis, Lehigh University, 2017.
- [40] Ugur Kuter, Dana S. Nau, Marco Pistore, and Paolo Traverso. A hierarchical task-network planner based on symbolic model checking. In *ICAPS*, 2005.

- [41] Gregg Rabideau, Steve A. Chien, and David McLaren. Tractable goal selection for embedded systems with oversubscribed resources. *JACIC*, 8(5):151– 169, 2011.
- [42] Sravya Kondrakunta and Michael T. Cox. Autonomous goal selection operations for agent-based architectures. In *IJCAI GRW*, Melbourne, Australia, 2017.
- [43] Nick A. M. Tinnemeier, Mehdi Dastani, and John-Jules Ch. Meyer. Goal Selection Strategies for Rational Agents, pages 54–70. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.
- [44] Mark A. Wilson, James McMahon, A. Wolek, David W. Aha, and B.H. Houston. Toward goal reasoning for autonomous underwater vehicles: Responding to unexpected agents. In 25th International Joint Conference on Artificial Intelligence (IJCAI) Workshop on Goal Reasoning, New York, NY, 2016.
- [45] M. Cashmore, M. Fox, D. Long, D. Magazzeni, and B. Ridder. Opportunistic planning in autonomous underwater missions. *IEEE Transactions* on Automation Science and Engineering, 15(2):519–530, April 2018.