Toward a Test Environment for Autonomous Controllers

Pablo Muñoz
Universidad de Alcala
Alcala de Henares, Madrid (Spain)

Amedeo Cesta, Andrea Orlandini
CNR – Italian National Research Council
ISTC, Rome (Italy)

Maria Dolores R-Moreno
Universidad de Alcala
Alcala de Henares, Madrid (Spain)

Abstract

In the last two decades, an increasing attention has been dedicated on the use of high level task planning in robotic control, aiming to deploy advanced robotics systems in challenging scenarios where a high autonomy degree is required. Nevertheless, an interesting open problem in the literature is the lack of a well defined methodology for approaching the design of deliberative systems and for fairly comparing different approaches to deliberation. This paper presents the general idea of an environment for facilitating knowledge engineering for autonomy and in particular to facilitate accurate experiments on planning and execution systems for robotics. It discusses features of the On-Ground Autonomy Test Environment (OGATE), a general testbench for interfacing deliberative modules. In particular we present features of an initial instance of such system built to support the GOAC robotic software.

Introduction

In the last two decades the ongoing work in robotics has been marked by the exponential growth of the available functionalities. This is a consequence of the advancements in both hardware and software technology. Still, robotic systems can conquer new locations and perform more complex tasks, also in environments which are unexplored.

As a result of deploying advanced robotic systems in unknown and dynamic environments, the control software required to achieve the mission goals shall deal with an important number of constraints. Thus, the advances in the Artificial Intelligence (AI) field seem to be naturally merged with the control of robotic systems in order to allow it to generate long term plans without (or little) human interaction. In this way, developments in planning and scheduling systems, such as task planning (Korf 1987; Fikes and Nilsson 1971), CSPs (Dechter and Pearl 1987; Do and Khambhampati 2000; Cesta, Fratini, and Oddi 2005), timelines (Muscettola 1994; Jonsson et al. 2000; Smith, Frank, and Jonsson 2000; Frank and Jonsson 2003; Fratini, Pecora, and Cesta 2008; Cesta et al. 2009; Chien et al. 2010) and, recently, the efforts interleaving planning and execution (Ambros-Ingerson and Steel 1998; Finzi, Ingrand, and Muscettola 2004; Py, Rajan, and McGann 2010) could be very valuable to control robots in dynamic environments with an increasing degree of autonomy.

These AI controllers are the top layer of complex tools that are usually made ad-hoc to control a specific robotic platform to perform determinate missions. Usually, to test and verify the correctness of an architecture, a small set of missions are carried out. Also, some parts of the control system could be evaluated in a standalone manner via particular test beds. But testing the robustness, adequacy and performance for the whole control architecture cannot be easily done; it requires to collect and to analyze relevant data from all the parts of the control system while the test bed covers more cases than the typical scenarios. This is currently an open and interesting problem in which not much work has been already done.

The paper presents a prototype of the On-Ground Autonomy Test Environment (OGATE)[1] as the initial result of an effort that aims to provide testing support while developing controller systems for space robotics missions. The OGATE environment would constitute an entry point to investigate these open problems in autonomous controllers as the combination of an engineering effort, identifying the requirements, designing and implementing a general environment to provide testing and verification tools for autonomous controller systems; and a research effort to discriminate the key factors in research on planning, scheduling and execution in order to evaluate the performance of autonomous controllers.

We also aim to provide support for ground segment facility in space robotics missions, hiding the complexity of the controlled system to the user. So, in the paper, the space robotics context is exploited as a real-world scenario, employed to test different solutions for deliberation and execution under the same conditions.

The paper is structured as follows: next section describes a space robotic scenario related to the GOAC project and exploited as a case study. Then, the objectives of the system are presented and, in the following, a brief description of OGATE and its functionality is described. The presentation of an initial deployment of the system and a description of what is a plug-in component for OGATE are given. And finally some conclusions are outlined.

[1]Funded by the ESA Networking and Partnering Initiative *Cooperative Systems for Autonomous Exploration Missions*.

A Space Robotic Case Study

Our interest in plan-based autonomy is also related to a recent participation in the Goal Oriented Autonomous Controller (GOAC) (Ceballos et al. 2011) project: an ESA effort to create a common platform for robotic software development. In particular, the GOAC effort combines several technologies: (a) a timeline-based deliberative layer which integrates a planner, called OMPS (Fratini, Pecora, and Cesta 2008), built on top of APSI-TRF to synthesize timelines and revise them according to execution needs, and an executive *a la* T-REX (Py, Rajan, and McGann 2010); (b) a functional layer (Bensalem et al. 2010) which combines a state of the art tool for developing functional modules of robotic systems ($G^{en}oM$) with a component based framework for implementing embedded real-time systems (BIP).

The GOAC system allows to implement controllers in a flexible way, i.e., for each robot or mission a different instance of the T-REX system can be deployed defining various cooperating reactors and their associated interactions, providing a scalable architecture. A T-REX agent is composed by a hierarchy of deliberative reactors. Each deliberative reactor has its own deliberative scope as well as planning latency and look-ahead, in charge of controlling a particular aspect of the mission and, interacting with other reactors sending goals and receiving observations. Then, a reactor exploits a planning system to generate plans and to monitor the execution following a sense-plan-act paradigm for goal oriented autonomy. It allows a divide and conquer approach in which the scope of each deliberative reactor could be refined by other more specific reactors. Firstly, the planning system employed in the T-REX deliberative reactors was the $EUROPA_2$ planning and scheduling framework (Jonsson et al. 2000; Bresina et al. 2005). For the GOAC project the planning system was replaced by the OMPS planning system, which exploit the APSI-TRF execution facilities.

Figure 1 represents a possible instance of the GOAC architecture. In the figure appears two deliberative reactors each one with its own planner and model over which to deliberate. These reactors are interconnected between them, and also, with a command dispatcher reactor. This one is in charge of sending commands to the functional layer (generally composed of different functional modules), while retrieving the observation and propagating them along the other reactors. With these data, the different deliberative reactors, using their respective deliberation models, could dynamically adapt their plans to the new circumstances.

Within the GOAC initiative, the DALA rover has been considered to simulate a robotic scenario as close as possible to a planetary exploration rover. DALA is one of the LAAS-CNRS robotic platforms that can be used for autonomous exploration experiments. In particular, it is an iRobot ATRV robot that provides a large number of sensors and effectors. It can use vision based navigation (such as the one used by the Mars Exploration Rovers Spirit and Opportunity), as well as indoor navigation based on a Sick laser range finder.

In this regard, DALA can be considered as a fair representative for a planetary rover equipped with a Pan-Tilt Unit (PTU), two stereo cameras (mounted on top of the PTU), a panoramic camera and a communication facility. The rover is able to autonomously navigate the environment, move the PTU, take high-resolution pictures and communicate images to a Remote Orbiter.

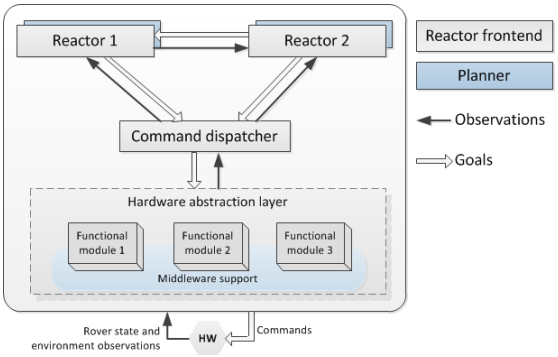

Figure 1: Representation of a GOAC instance

The considered mission goal is a list of required pictures to be taken in different locations with an associated PTU configuration, and to communicate them to an Orbiter when it is visible to the robot. Also, the rover must operate following some operative rules to maintain safe and effective configurations (the reader may refer to (Ceballos et al. 2011) for further details). To deal with the objectives and the operative rules, the OMPS deliberative is in charge of synthesizing a sequence of actions that, starting from the environment state, robot state and goals, reach a final state in which the goals are satisfied.

A possible mission actions sequence is the following: navigate to one of the requested locations, move the PTU pointing at the requested direction, take a picture, then, communicate the image to the orbiter during the next available visibility window, put back the PTU in the safe position and, finally, move to the following requested location. Once all the locations have been visited and all the pictures have been communicated, the mission is considered successfully completed.

Within the GOAC project was also exploited the ExoMars rover model using the ESA 3DROV simulator suite (Poulakis et al. 2008) that allows early-stage virtual modelling of terrain and mobile robots systems. The system is composed of multiple modules connected through standardized interfaces, being the most important the *Simulation Framework*, that is the ESA's Simsat, responsible for the execution and scheduling of the simulation and the *Generic Controller* that manages the onboard flight software to enable to connect software modules to control the rover. Also, it includes the *Environment block* in charge of the timekeeping, terrain and atmospheric conditions, and the *Visualization Environment*, a front-end that provides real-time visualization of the simulation progress.

On-ground support for autonomous control

The current investigation is triggered in the space environment where the use of software for on-board autonomy is often perceived as loosing control on critical mission com-

ponents (namely a space robots, a spacecraft, etc.). For sure current complexity of software for autonomy is quite high and such a complexity reinforce the general skepticism toward its wide use in the space environment. To cope with the problem we have conceived the idea of creating a software environment to be used on-ground to facilitate the demonstration and testing of software for autonomy. Such an environment can also represent the seed for a future knowledge engineering environment for autonomous controllers. Indeed the first goals for such an environment are: (a) facilitating the use of autonomous controllers; (b) allowing the use of different solution for autonomous control (e.g., toward a plug-and-play style in their use); (c) enabling the comparison of different solutions gathering reliable execution data on a given mission. For the time being we are also making an additional assumption: we are focusing our attention toward the deliberative part of the autonomous control the part that can be referred to as the one performing "planning and execution" hence we assume that the physical system is accessible through a functional interface (in GOAC, a-la G^{en}oM (Mallet et al. 2010)) or through a robotic operating system.

To make our general goal clear let us refer to the space robot domain introduced above. Our research plan is to develop an easy-to-use system able to (i) interface different planning and execution solutions with a same robot (or its simulator) and (ii) to automatically generate realistic test bed scenarios presenting an increasing complexity.

Here, we consider missions related to the space robotics case study. Those missions consist in using a determinate autonomous control architecture over a robotic platform, return some science objectives (i.e., scientific pictures) taking also into account a set of constraints such as the availability time windows for Remote Orbiter communication, or the time period in which science targets are available (some events may be time bounded). The goal is to progressively increase the difficulty of the missions, aiming to stress the planning and execution system and, also, to collect performance information exploiting a real robotic platform like DALA as well as a simulator suites such as 3DROV (Poulakis et al. 2008).

Afterward, we are planning to investigate how different ways of conceiving the planning and execution task can be compared. In particular, we need to compare them using the same functional support, as fig. 2 shows, and identical missions.

Although we can find stress tests for planning systems (for example the International Planning Competition (IPC)[2] in which PDDL-based planners (Gerevini and Long 2005) performance is evaluated based on quantitative criteria), they are standalone tests in invariant conditions. This means that there is no interference due to a dynamic environment (e.g., climate changes, external agents, new goal opportunities, etc.) or to changes in the system (e.g., malfunctions or failures) during a mission that affects directly the planning and execution system. Also, newest technologies interleaving

[2]The IPC is usually co-located with the International Conference of Automated Planning and Scheduling (ICAPS).

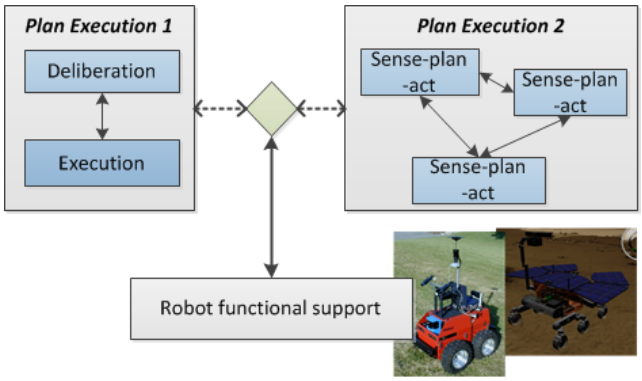

Figure 2: Different configurations for planning and execution controlling the same robotic platform

planning and execution may use different schemas, for instance, in terms of number of deliberative components or their scope and specialization. Then, a solution plan can also be found by different cooperating autonomous systems. This entails also the need of investigating how these schemas potentially affect the planning process. And, at the best of our knowledge, a methodology to compare performance metrics in realistic scenarios, such as the space robotic mission domain considering not-nominal conditions, is still missing. So, this constitute an interesting open research issue.

The OGATE infrastructure

We can briefly define a mission as a set of goals that are solvable by an autonomous agent. This autonomous agent is composed of two parts: i) the platform, which could be a simulator or a robot, and, ii) the control architecture to manage the platform, a set of software components that work together to accomplish the objectives defined in the problem. So, considering an autonomous agent, one or more of these components will be a deliberative component, which are responsible of managing the long term planning to accomplish the objectives of the mission.

Currently, there exists a high number of technologies available to define problems for autonomous agents and their corresponding model of the world (usually called domain), which produce high level plans, which must be decomposed in order to be executed by a robotic platform that accepts low level commands. To deal with the complexity of operating the platform and controlling the execution and decomposition of the high level plans, the control architectures are typically structured in three levels (Gat 1998): i) the deliberative layer formed by one or more deliberative components in charge of the long term planning and scheduling ii) the low level support, called functional level that controls the platform using low level commands, and iii) the executive layer, an intermediate level that decomposes the actions produced by the deliberative layer into commands accepted by the functional level. Newer developments interleaves the two upper layers (deliberative and executive) into a decisional layer, in which planning and execution are highly coupled.

So, if we are interested in evaluating the performance of a control architecture, we need to take into account all the components, not only the deliberative components. The configuration of the architecture, that is, the hierarchy built on top of a set of different components and how they are connected, plays a fundamental role: some planning technologies generate a complete plan before execute it, while others generate partial plans, interleaving planning and execution in a loop. This implies questions such as the different delays interchanging data between layers that affects the performance of the system. To cover that issue we take the functional support for the robotic platform to control as an invariant part of the system, while employ different technologies for planning and execution over the same domain.

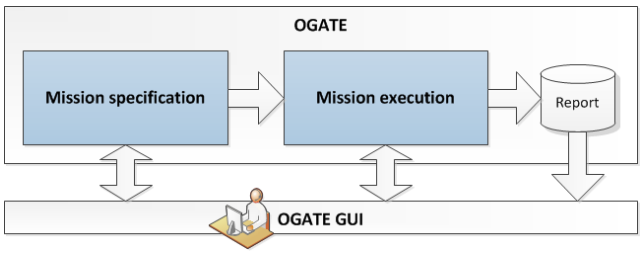

Figure 3: The OGATE environment

To deal with these questions, OGATE aims at providing an environment to test features of goal oriented controllers and to obtain quantitative comparison based on accurate experiments and also qualitative analysis allowed by inspection and visualization of software internal monitors of the controlled system. The infrastructure of the system relies on three main modules as seen in fig. 3 to support a general test bed environment for autonomous agents:

– **Mission specification:** to define the functionality and goals, first it is required to specify the configuration of the mission: the components of the control architecture and the platform over which they operate. In this way, the system provides a convenient mode to allow the user to configure some components of the controlled system, such as the deliberators. Typically, these AI controllers work with a domain and a problem. First one defines the interactions between different elements of the world, and the last one includes some initial facts and the desired objectives of the mission. The specification of a mission testbench could consist of an evaluation of different control architectures or various configurations for a control architecture to select the best one for a particular mission, or the evaluation of the performance for a particular control architecture over a set of missions, to evaluate and improve the components employed. The OGATE system will be able to support these tests in an automated way.

– **Mission execution:** the mission specification includes the configuration of the different components (executives, deliberatives, etc.) involved in the mission. The execution support of OGATE provides a framework to deal with the complexity of the underlying architecture of the different components, to execute the user defined mission and to gather the relevant data in order to give it to the user. Also, during the execution, the user must be able to interact with the controlled system, modifying internal parameters or including new mission goals to change the nominal execution, in order to test the robustness of the system and the replanning capabilities of the planning and execution components, or to include new science opportunities for real missions to maximize the science return. So, using specific defined interfaces, OGATE is able to access to the different components to control and gather the relevant data.

– **Report:** from the previous data, a research effort to identify and to develop useful metrics for comparing the performance of different deliberative layers will be addressed, in order to obtain strong conclusions when comparing different deliberative components under the same conditions. OGATE will provide a human-legible report of the mission when the required test have been performed.

A first contact with the environment

For the first deployment of the OGATE environment we employ the GOAC architecture. The mission specification relies on the OGATE infrastructure to provide the configuration of the mission by the user, that is, select domain and problem, and the deliberative components that will be attached to the GOAC architecture. Also, it will be possible to define different platforms to control, with their respective functional layers. The configuration of the different reactors for the architecture (i.e. command dispatcher or functional layer) will be easily set by the user using a provided graphical interface, encapsulating the complexity of dealing with the underlying GOAC architecture. The mission specification generates a configuration file as shown in fig. 4 to be used by the mission execution.

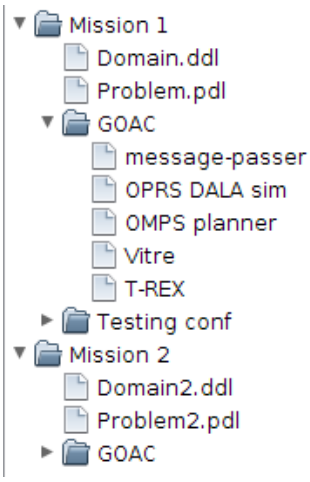

Figure 4: OGATE mission specification

The mission execution takes the configuration defined by the user and it attaches the different components to the OGATE system to execute them in a coordinated way.

Some of these components are basic components over which OGATE has no control, it only executes them because they are required (for example, a simulation platform will be executed before execute the functional layer that controls the platform), while other components will be controlled directly by OGATE, such as the deliberative component. These components will be defined as OGATE plug-in components (more details in next section). For deliberative plug-in components a GUI allows the user to include new goals or to modify the internal state during the test, showing the relevant data in real-time.

From the data gathered by the plug-in components and other relevant data generated by the system, OGATE will provide a report to measure the performance for these components and for the whole configuration.

At this moment we are focused on the deliberative capabilities, and for the initial test we are interested in two points:

- Starting from the timelines approach and with the APSI-TRF support, deploy an automated problem generator and to perform some initial tests to evaluate the planner, increasing the problem hardness. This generator must be based on critical factors that affect directly to the goal oriented controllers, and the problems must be defined using a set of parameters which denote the problem hardness with valid and objective criteria.
- When the system is mature, we plan to interchange the configuration and technology employed for the deliberative and executive process, using different instances of GOAC for interleaving planning and execution, and other solutions based in different paradigms, such as the *plan-then-act schema*, like MOBAR architecture does (Muñoz and R-Moreno 2013).

The deliberative plug-in

A OGATE plug-in is the one that implements some functionality, giving some access to OGATE to control its execution, or to retrieve data from it through specific interfaces. Focusing on the deliberative component, the OGATE mission execution module should access and control this component, then its functionality (the planning service) must be encapsulated in a OGATE plug-in that could have three connections with the OGATE system as fig. 5 shows. In the figure the deliberative component functionality is accessed by the OGATE system through the following interfaces:

- **Control interface:** provides the basic functionality to run, pause or stop the component safely.
- **Data interface:** supplies a bidirectional channel to retrieve the relevant data and to modify the internal state of the component. For deliberative components it also provides a function to include new goals during the execution.
- **Component GUI:** it is possible to include a specific user interface for that component in the OGATE GUI.

In order to implement the different interfaces, a set of Java methods is given in a template for every interface. For the first and the second ones, the methods are calling such as `boolean stop()` or `float getFloatVal(String valId)`. As it is possible that the component is developed in C/C++, OGATE will also include some functionality to connect these components giving some standard code to made the interconnection between OGATE and the component. The implementation of the functionality of these methods are responsibility of the component developer. For the component GUI, OGATE can integrate a Java GUI inside its environment with a minimum modification of the code; graphical components developed in other technologies must be executed outside the environment or executed by OGATE as a basic component.

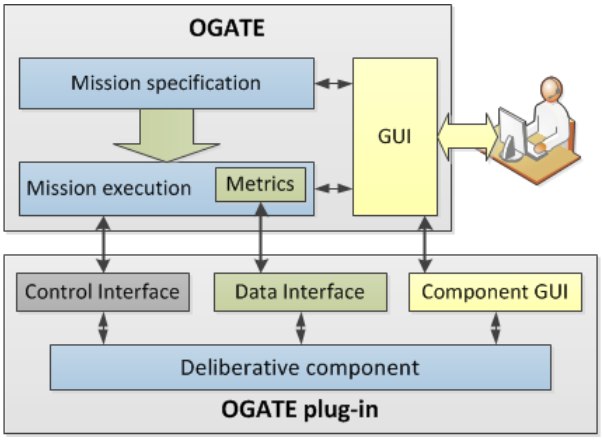

Figure 5: A deliberative component plug-in

From the different plug-ins attached to the OGATE system, the information gathered through the data interface will be passed to the Metrics module, which takes the relevant data and generates a report based on specialized metrics. This is the final result of the execution of the OGATE system.

Current deploy of OGATE

The deployment of an automated control architecture is usually a complex task, that requires some knowledge about the different components: how are they interconnected, which one must be executed first to be coupled with the others, etc.. So, if we are interested in disseminating our work, we must spend some time in preparing manuals and help new users.

For these new users, when they are able to start the system, a huge amount of data is (generally) generated and presented via the command line or via specialized graphical interface. But depending on the users preferences, they may want to pay attention to some particular data.

The current state of the OGATE system integrates Knowledge Engineering capabilities, using the different components that forms the GOAC architecture, what allows us to:

- Create new problems of the space robotic case study domain (mentioned in previous section) using an automatic generator that defines the hardness in function of the number of science tasks and the time in which the Remote Orbiter is visible.

– Create scenarios through a GUI to specify the planning and execution configuration used to solve the problem(s) defined. Currently, the T-REX configuration is fixed to one deliberative reactor, the command dispatcher reactor (connected to the functional support of the robot) and a visualization tool provided by the T-REX framework (it visualizes the current state of the different timelines contained in the domain). So, the user only can change the planning service employed by the deliberative reactor and the domain and problem to resolve.
– Define components that will be run without the OGATE control such as the functional support of the robotic platform or the simulator. Actually in our implementation, the DALA robot simulator is only available.
– Integrate different plug-ins; OGATE not only accepts deliberative plug-ins, other components could be attached to the system, such as the visualization tool of T-REX.
– Attach the different components to generate an instance of the system, which will be executed by OGATE without any intervention of the user.
– Focus on what we are really interesting since inside OGATE is possible to define which components generate relevant data, and then, show only that information inside a unique GUI.

The cycle for the execution of a control system within OGATE could be seen in fig. 6. Starting form the different components available (the T-REX engine, a timeline-based planner, domain and problem, and a robotic platform), the user can exploit the mission specification module of OGATE to generate a configuration file that defines what are the problems to address, and with which instance of T-REX the problem will be solved within the mission specification GUI (shown at the top left window). After that, the mission execution module takes the user choice and creates an instance of T-REX with the selected planner and the functional support, and it attaches them correctly. In fig. 6 is possible to see (bottom right) OGATE running the current available instance of T-REX with the OMPS planner as the deliberative component inside an OGATE plug-in. This plug-in implements the control and data interface to allow OGATE to control the planning service and to retrieve the relevant data produced. Also, there is another OGATE plug-in to include the T-REX' visualization tool inside the OGATE GUI. The functional support is the DALA simulator that requires specific components to be executed without being controlled by the OGATE system. One of them is the `mp-oprs`, a message passer service for communication with the robotic platform exploited by OGATE, for which output is presented inside the OGATE GUI. Also, OGATE maintains a registry or log for the different components, which could be valuable for debugging the controlled system. Finally, OGATE must present the relevant data gathered to the user.

Although OGATE provides some advantages, we are currently working on expanding its capabilities for making execution more "user-comprehensible". At this moment OGATE shows the information of the components, but not the interaction between them. This is an important issue that has to be addressed to better understand what and how the system works. Also, through the GUI, our intention is to provide a framework that allows the user to support different planning and execution systems via a plug-in style, trying to ease the deployment of complex systems. The support is not only applicable to the deliberative layer, but also to the functional support, in order to exploit different robotics platforms or simulators suites. Finally, our intention is to investigate the metrics to define the performance for planning and execution systems, implementing objective and comprehensive comparison reports for autonomous controllers inside the OGATE system.

Conclusions

Control architectures for autonomous robots are complex software systems, not only from a design and implementation perspective but also from a research point of view since we want to analyze and (possibly) improve their performance. This is in part a consequence of the lack of a testbench that allows the user to define metrics and obtain results from the execution of different tests (by inspection of the relevant parameters).

Also, it is usually very complex to deploy and execute an autonomous control architecture, which generally is formed by different interconnected components. Besides, once the systems is executing, the operator gets a huge amount of non human-legible information, confusing the user and keeping him/her away from the relevant data.

For this reason we are working on the OGATE system, which aims to address part of this gap, providing a general testbench to allow easy deployment of an autonomous architecture: giving the OGATE plug-ins and configurations files, it will be easy to deploy any system. It will be only required to load the specific files and run it. Then, specialized graphical interfaces will be enabled to show the relevant data to the user, encapsulating the complexity of the underlying functionality. Finally, OGATE will also offer an automated testbench to obtain quantitative comparison based on accurate experiments, which leads to human-legible reports with well defined metrics.

Acknowledgements

Pablo Muñoz is supported by the European Space Agency (ESA) under the Networking and Partnering Initiative (NPI) *Cooperative systems for autonomous exploration missions*. CNR authors are partially supported by the Italian Ministry for University and Research (MIUR) and CNR under the GECKO Project (Progetto Bandiera "La Fabbrica del Futuro"). Authors want to thank to the ESA's technical officer Mr. Michel Van Winnendael for his continuous support.

References

Ambros-Ingerson, J., and Steel, S. 1998. Integrating Planning, Execution and Monitoring. In *AAAI*, 83–88. AAAI Press.

Bensalem, S.; de Silva, L.; Gallien, M.; Ingrand, F.; and Yan, R. 2010. "Rock Solid" Software: A Verifiable and Correct-by-Construction Controller for Rover and Spacecraft Func-

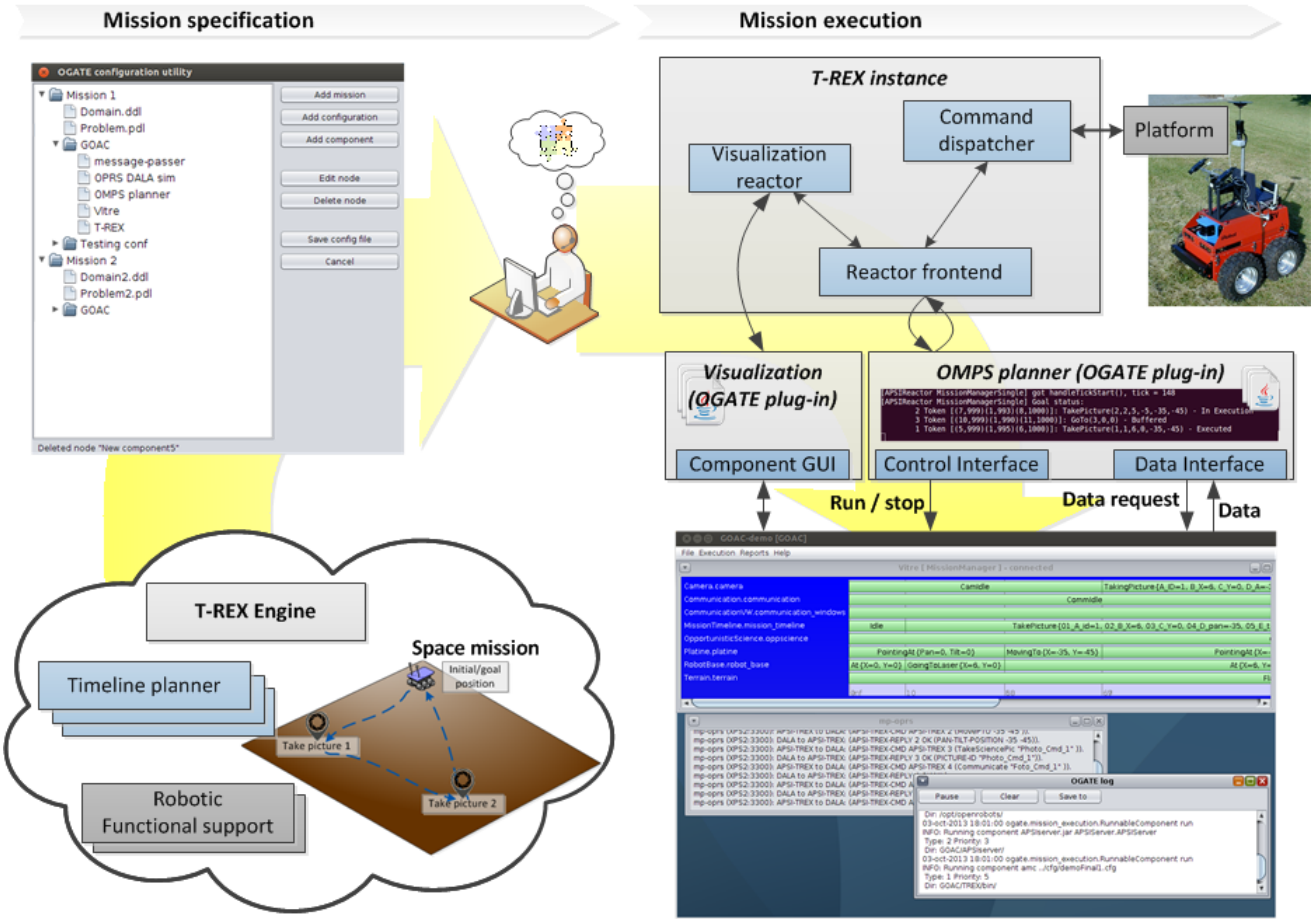

Figure 6: A first deploy of the OGATE system

tional Levels. In *i-SAIRAS-10. Proc. of the 10th Int. Symp. on Artificial Intelligence, Robotics and Automation in Space*.

Bresina, J.; Jonsson, A.; Morris, P.; and Rajan, K. 2005. Activity planning for the Mars Exploration Rovers. In *in Proc. of the 15th International Conference on Automated Planning and Scheduling*.

Ceballos, A.; Bensalem, S.; Cesta, A.; Silva, L. D.; Fratini, S.; Ingrand, F.; Ocón, J.; Orlandini, A.; Py, F.; Rajan, K.; Rasconi, R.; and Winnendael, M. V. 2011. A Goal-Oriented Autonomous Controller for Space Exploration. In *ASTRA 2011 - 11th Symposium on Advanced Space Technologies in Robotics and Automation*.

Cesta, A.; Cortellessa, G.; Fratini, S.; and Oddi, A. 2009. Developing an end-to-end planning application from a timeline representation framework. In *IAAI-09. Proc. of the The Twenty-First Innovative Applications of Artificial Intelligence Conference*.

Cesta, A.; Fratini, S.; and Oddi, A. 2005. Planning with Concurrency, Time and Resources: A CSP-Based Approach. In Vlahavas, I., and Vrakas, D., eds., *Intelligent Techniques for Planning*. Idea Group Pubhishing. 259–295.

Chien, S.; Tran, D.; Rabideau, G.; Schaffer, S.; Mandl, D.; and Frye, S. 2010. Timeline-based space operations scheduling with external constraints. In *ICAPS-10. Proc. of the Twentieth International Conference on Automated Planning and Scheduling*.

Dechter, R., and Pearl, J. 1987. Network-based Heuristics for Constraint-Satisfaction Problems. *Artificial Intelligence* 34(1):1–38.

Do, M. B., and Khambhampati, S. 2000. Solving Planning-Graph by Compiling It Into CSP. In *ICAPS-00. The Fifth International Conference on Artifcial Intelligence Planning and Scheduling*, 82–91.

Fikes, R. E., and Nilsson, N. J. 1971. STRIPS: A New Approach to the Application of Theorem-Proving to Problem-Solving. *Artificial Intelligence* 2(3):189–208.

Finzi, A.; Ingrand, F.; and Muscettola, N. 2004. Model-based executive control through reactive planning for autonomous rovers. In *In Proc. of 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems*.

Frank, J., and Jonsson, A. 2003. Constraint-based attribute and interval planning. *Journal of Constraints* 8(4):339–364.

Fratini, S.; Pecora, F.; and Cesta, A. 2008. Unifying Plan-

ning and Scheduling as Timelines in a Component-Based Perspective. *Archives of Control Sciences* 18(2):231–271.

Gat, E. 1998. Three-layer architectures. In Kortenkamp, D.; Bonasso, R.; and Murphy, R., eds., *Mobile Robots and Artificial Intelligence*. AAAI Press. 195–210.

Gerevini, A., and Long, D. 2005. Plan constraints and preferences in PDDL3. In *Proc. of the Fifth International Planning Competition*.

Jonsson, A. K.; Morris, P. H.; Muscettola, N.; Rajan, K.; and Smith, B. D. 2000. Planning in interplanetary space: Theory and practice. In *the 5th International Conference on Artificial Intelligence Planning Systems*.

Korf, R. E. 1987. Planning as Search: A Quantitative Approach. *Artificial Intelligence* 33:65–88.

Mallet, A.; Pasteur, C.; Herrb, M.; Lemaignan, S.; and Ingrand, F. 2010. GenoM3: Building middleware-independent robotic components. In *2010 IEEE Proc. of the International Conference on Robotics and Automation*.

Muñoz, P., and R-Moreno, M. D. 2013. Model-Based Architecture on the ESA 3DROV simulator. In *In Proc. of the 23rd ICAPS Application Showcase*.

Muscettola, N. 1994. HSTS: Integrating Planning and Scheduling. In Zweben, M. and Fox, M.S., ed., *Intelligent Scheduling*. Morgan Kauffmann.

Poulakis, P.; Joudrier, L.; Wailliez, S.; and Kapellos, K. 2008. 3DROV: A planetary rover system design, simulation and verification tool. In *International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*.

Py, F.; Rajan, K.; and McGann, C. 2010. A Systematic Agent Framework for Situated Autonomous Systems. In *AAMAS-10. Proc. of the 9th Int. Conf. on Autonomous Agents and Multiagent Systems*.

Smith, D.; Frank, J.; and Jonsson, A. 2000. Bridging the gap between planning and scheduling. *Knowledge Engieneering Review* 15(1):47–83.