

# The Remote Agent Executive: Capabilities to Support Integrated Robotic Agents

Barney Pell<sup>§</sup>   Gregory A. Dorais<sup>†</sup>   Christian Plaunt<sup>†</sup>   Richard Washington<sup>†</sup>  
pell | gadorais | plaunt | richw @ptolemy.arc.nasa.gov  
NASA Ames Research Center, Mail Stop 269-2  
Moffett Field, CA 94035-1000

## Abstract

The Remote Agent (RA) integrates a broad spectrum of robotic activities, including planning, scheduling, execution, monitoring, failure detection, diagnosis, and recovery. The RA Executive (EXEC) can be viewed as the core of the agent. EXEC enables software developers to think about the robot at a higher level; it also supports the reuse of knowledge and code across multiple robot applications. EXEC's capabilities include a high-level procedural action-definition language, services for resource management and configuration management, support for executing flexible closed-loop plans and command sequences, support for replanning, and a framework for specifying fault responses including safe modes and responses to plan failures.

We believe that these capabilities are required for most autonomous agents, and that executives and execution capabilities will become more essential as we attempt to develop autonomous agents of increasing capability. Moreover, we deem modularity, variable autonomy, robustness, and support for tools and testing to be general architectural properties necessary for design, deployment, and operation of executives within a complex mission. To this end, we are remodularizing the executive by providing each capability as a separate component so that the individual capabilities can be standardized and shared across different agent architectures.

## Introduction

An executive coordinates the runtime activity among a set of modules involved in a control system. Such coordination is necessary to synchronize coordinated activities, avoid resource conflicts, enforce operational constraints, and ensure safety in the presence of failures.

The capability level of an executive has direct impact on the design and complexity of the entire system. If there is no executive in the system, the individual modules must be commanded separately. Also, all the coordination requirements must be performed open-loop or the modules must know about each other

and coordinate themselves directly. This increases the complexity of both the external commanding system and the functional modules themselves.

Most deployed spacecraft (s/c) control systems today have a simple sequencing engine on-board, which accepts a timed sequence of commands and issues those commands to the functional modules. The current systems do not support context-sensitive or conditional activity, or decomposition of high-level activities into lower-level commands. One result is that the duration of activities must be determined open-loop by the ground system. These limitations cause significant complexity in the ground system, as it must reason about every detail and every interaction of the control system. Even for ground systems that employ automated planners to ease the burden, the level of detail required for planning down to the lowest-level activities makes the planning problem computationally difficult. The low-level detail also makes the knowledge encoded and used by the ground system (including the planner) cumbersome to extract and validate, and error prone.

Moreover, since current executives do not understand the causal structure of the sequences they execute and do not perform context-sensitive activities, the resulting behavior is brittle. The sequence tends to break if anything goes wrong, causing loss of mission goals and missed opportunities. Missions that cannot afford such losses at critical periods wind up pushing further complexities into the functional module software, and onto the fault protection software. For example, in the "sequence rollback" approach, the functional modules that receive the low-level commands must be able to execute the commands differently depending on whether they succeeded or failed last time, so that the same sequence can be re-used in the case of failures. This approach leads to software that is almost incomprehensible, unmaintainable, and certainly not reusable.

## The Remote Agent

To address these control system limitations, we have developed the Remote Agent (RA). RA is designed to control s/c autonomously for extended periods of time, and is being demonstrated on the Deep Space One s/c (DS-1) scheduled to launch July, 1998 (Bernard *et al.*

<sup>§</sup>Research Institute for Advanced Computer Science

<sup>†</sup>Caelum Research Corporation

1998). A prototype of RA was also tested on a complex simulated Saturn Orbit Insertion (SOI) scenario for the Cassini s/c (Pell *et al.* 1998a). In this paper we use s/c to refer to any extraterrestrial craft including rovers. However, RA is not limited to controlling s/c.

RA is composed of the planner-scheduler (PS), the model-based mode identification and recovery system (MIR) for fault diagnosis and repair, and the executive (EXEC) (Pell *et al.* 1998a). This paper focuses on EXEC.

## The Remote Agent Executive

EXEC is event-driven and goal-oriented and forms the core of RA. It provides a language and a framework in which software designers can express how planning, control, diagnosis, and reconfiguration capabilities are to be integrated into an autonomous system. It can request and execute plans involving concurrent activities that may be interdependent, where the success, timing, and outcomes of these activities may be uncertain. It provides a language for expressing goal-decompositions and resource interactions. When interpreting this language at run time, the executive automates the decomposition of goals into smaller activities that can be executed concurrently. This automates aspects of the labor-intensive sequencing function in spacecraft operations and raises the level of abstraction at which the ground system or onboard planner must reason. The framework also supports a close integration between activity decomposition and fault responses. This leads to more robust execution, avoids loss of mission objectives, improves mission reliability and resource utilization, and simplifies the design of the entire software system.

As an integrating technology, an executive enables each software component to be expressed independently, but controls the interactions among those components so that the overall activity of the system achieves global constraints as dictated by a high-level plan. This increase in modular software design enables reduced development and integration costs and increased software reuse both within and across missions. This capability also enables a new generation of adaptive controllers with multiple control modes. A smart executive switches control modes of these controllers based on feedback from the environment.

Our understanding of the role and requirements of executive capabilities has evolved since we began our prototyping effort on the SOI scenario. That work built on executive components the Cassini engineers had developed, including a configuration manager, resource manager, and fault protection framework (Brown, Bernard, & Rasmussen 1995) and extended these capabilities to support an on-board planner and model-based diagnosis and repair system. On DS-1, we moved from a prototype to an operational flight software system that was part of a much larger project. This experience provided lessons about ad-

ditional requirements that were critical to successful deployment in an actual mission (Aljabri *et al.* 1998). Our current work revolves around extending the capabilities of the individual components and addressing critical features in response to the lessons from our flight experience.

The rest of this paper is structured as follows. The next two sections describe the EXEC components and capabilities we think are general to most autonomous control systems and then general architectural properties necessary for design, deployment, and operation within a complex mission. We then discuss related work and future directions.

## EXEC Components

EXEC has the following capabilities which we consider useful for executives in general:

- Flexible plan execution
- Configuration management
- Resource management
- Action-definition language
- System-level fault-protection support

While these capabilities are tightly coupled within EXEC, we will describe each as a separate component with independent functionality. We are currently redesigning EXEC to reflect this modular architecture, as discussed later.

### Flexible Plan Execution

One major capability of EXEC is to support execution of plans generated by PS. Several features of plan execution are important in our applications.

An important set of features concerns the plan representation itself. PS generates *flexible, concurrent, temporal plans* (Pell *et al.* 1998a; Muscettola *et al.* 1995). This is important for control systems in which many subsystems conduct concurrent activity, and in which execution must be robust in the face of uncertainty. Despite the flexibility in the plans, it is important that EXEC execute the plans correctly and efficiently. EXEC's own processing of the plan introduces some latency during execution. This introduces issues of temporal granularity and propagation times into the system design (Muscettola *et al.* 1998).

Planning can be an infrequent activity due to resource and informational limitations. Hence, EXEC supports the execution of infrequently-generated plans (Pell *et al.* 1997). To avoid interrupting activities at the boundary between execution of one plan and the next, EXEC supports *smooth plan transitions*, in which the same activities can be referred to between subsequent plans. Thus, a new plan can be smoothly merged into the previous execution context.

EXEC supports *concurrent planning and execution*, in which a new plan can be generated while a previous plan is executing (Pell *et al.* 1997). This requires

the planner to make use of current execution information. Also, EXEC must track the assumptions under which the new plan is being generated, and abort the planning process if those assumptions become invalid before the new plan is executed.

Monitoring the execution and aborting (under conditions of failure or by command) is also an important requirement. When the execution of a plan is aborted, EXEC terminates activities executing under the current plan (including any replanning activity currently underway) and invokes a procedure to establish a stable configuration, called a *standby mode*, from which a new round of planning can be initiated (Pell *et al.* 1997).

EXEC supports a mechanism by which users can switch between closed-loop plan execution, in which EXEC requests plans based on current state information, and open-loop execution, in which EXEC runs plans possibly generated by mission operators and sent to the s/c (Pell *et al.* 1998c). For this feature to be useful, it is important to support smooth plan transitions in the open-loop context as well.

In addition to the ability to execute high-level plans, it is also important to support the functionality of traditional sequencers. While such sequences are less flexible (and thus less robust in the face of uncertainty), they can be easier to analyze and predict, and hence more comfortable for mission operators. EXEC supports the ability to run traditional sequences, and also to embed sequences within flexible plans. This enables a form of variable autonomy (discussed later) in which operators can choose to control some aspects of operation via high-level goals managed by the agent but control other aspects by low-level commanding. Even when EXEC is running sequences, it gives operators access to a rich language in which they can express time-driven, event-driven, and context-sensitive relationships between activities.

## Services for Configuration Management

A challenge in designing software for a complex system, such as a s/c, is that the system may be in a wide variety of configurations. The physical components of a s/c, their topology, and the state of each component define a s/c configuration. At hardware design time, devices may be added or removed from the system (for example, a new star tracker), and details of individual hardware components may change. At run time, devices may be turned on or turned off, healthy or malfunctioning, and in a number of possible operational states.

EXEC's configuration manager (CM) insulates the software for each of the many activities in a complex system from the details of the configurations that those activities require. If the configuration details change, whether at run time or design time, the activity code can remain unchanged.

CM maps configuration requirements to code that

establishes configurations satisfying those requirements. The code that establishes configurations must reason about inter-dependencies among devices, and it must adapt to hardware changes. To simplify development of such code, we have designed a set of generic configuration management procedures, in which the user provides hardware details (including connectivity diagrams) and instantiates the generic procedures.

Once a configuration satisfying the required properties is established, CM tracks this configuration and signals an error if these properties are violated. In such failure conditions, CM interacts with the fault-protection capabilities to establish a potentially new configuration so that the code requiring it can continue to operate correctly.

In EXEC, individual device knowledge is implemented based on a library of generic device management routines. Devices and classes are formalized using generic descriptions. Individual devices, switches, etc., are then modeled as instances of these classes. Consider the following example:

```
(define-device-class :camera
  :power-function #'fsc-power-request
  :talk-function #'camera-talk-msg)

(define-device :camera_A :camera
  :powered-thru :power_bus_1
  :switched-thru :fsc_camera_sw1
  :ready-state ( (:health_state :ok)
                 (:power_state :on)))
```

Based on these device idioms, EXEC has generic procedures defined for device configuration and management. With a single construct, CM will select a device of a given class, achieve its ready-state, and then lock the properties of the ready-state and maintain them during the execution of an activity. Note that achieving a state of one device may require readying other devices it depends on, recursively.

Based on the camera definition above, the construct

```
(with-selected-device :camera
  (take-pictures))
```

would select a camera (say *camera\_A*), achieve its ready-state of being powered on and healthy, and then take pictures within a context that ensures that the health and power of the camera are maintained throughout picture taking.

## Services for Resource Management

Scarce resources are a fact of life on s/c. On-board processes contend for a limited number of devices. Shared resources, such as energy and data storage, are subject to hard (and relatively severe) limits, as well as environmental influences (such as solar exposure for battery charging) that change these limits over time.

Currently, s/c control sequences use worst-case estimates of resource utilization. However, the uncer-

tainty in the environment and its effect on the true resource utilization can make estimates inaccurate. For example, in the rover domain, navigation over a fixed distance can use widely varying amounts of energy depending on the terrain the temperature. A loose worst-case estimate can lead to under-utilization of the *s/c*; an underestimate can lead to aborts in plan execution.

A *s/c* resource manager must be able to manage and coordinate utilization of these scarce resources. First, it must provide a basic locking mechanism that maintains devices in particular states throughout the time interval for which the device state is needed by a process. Second, it must handle continuous and varying resources by responding to immediate resource requests from tasks, tracking resources over time, and allowing resources to be reserved for future tasks.

EXEC's *property-lock manager* manages a set of concurrent control tasks. Each control task requires a set of *properties* to be established and maintained over some period of time, e.g., taking pictures with a camera requires that the camera be on and functional. If some other activity requires the camera to be off, these two activities compete for the resource of controlling the camera's power state. EXEC achieves, maintains, and monitors properties required for each task, and resolves task-resource conflicts. Of course, these properties can change despite the best efforts of the software system to maintain them, e.g., switches on a *s/c* sometimes change state accidentally. The underlying fault-protection support constructs handle property violations.

Since property achievement may not be instantaneous, the resource manager also records the fact that the property has been requested and is already being worked on. This prevents another task from attempting to achieve the same property.

The EXEC resource manager is being extended to manage more complex resource requirements. The resources are not all-or-none, as in the current resource manager, but allow partial resource allocation (as with energy or data storage). In addition, the resources may be reserved for future use, so that future tasks may be scheduled.

The primary functions of the resource manager being developed are to detect changes in resource availability, adapt the resource allocation to respond to the resource changes, and manage resource requests and queries from tasks. The resource availability may change for a variety of reasons: degraded components (battery use), environmentally affected performance (solar cells), or other causes. The resource manager will have multiple strategies for adapting to resource changes and conflicts. These strategies will be selectable remotely and include task abortion (the simplest and most severe), shedding of low-priority tasks, and a novel idea of temporarily borrowing resources from other tasks.

## High-Level Procedural Action-Definition Language

*S/c* control software should provide a concise way of specifying the control and behavior of the various *s/c* devices and the core components of EXEC itself, and insulate the system developer from low-level interactions with hardware and software.

EXEC is built upon a powerful, extensible, high-level language, called Execution Support Language (ESL) (Gat 1996). This language was developed in response to the needs of developing EXEC and the systems it controls. ESL includes the following features:

- A top-level interpreter (shell)
- Constructs for managing task networks and concurrent threads of execution
- Conditional task execution
- Robust exception handling with constructs for condition monitoring and maintenance, multiple recoveries, and cleanups
- Event signaling, response, and call-backs
- A deductive memory database
- Constructs for relative and absolute timing

A language feature we found important is *language extension*. This enables us to define new constructs which encapsulate procedures we used repeatedly and to enforce programming conventions and architectural constraints. The lack of this feature motivated us to develop ESL as a successor to RAPS (Firby 1978), which we used during our prototype effort (Pell *et al.* 1998a).

## System-level Fault Protection

Fault protection enables a system to recover from anomalous conditions such as obstacles to navigation or device malfunctions. *S/c* engineers distinguish *component-level* fault protection, which tries to restore functionality of an individual subsystem, from *system-level* fault protection. In the latter, the overall system must be protected even if the component that failed is not identified, isolated, or recovered.

Writing fault-protection software is a significant part of a *s/c* software engineering effort because there are many subsystems and interactions to consider. To handle fault protection across subsystems, many constructs are needed. These include events and responses, timers, the ability to achieve standby mode, property tracking, constrained recoveries, and waiting for confirmation before taking action.

Our executive fault-protection component provides a robust layer for fault monitoring, detection and recovery. EXEC's fault-protection system provides:

- An event-response architecture, which has multiple threads of control listening for events (e.g., under-voltage trips or extended loss of communication),

which execute well-defined responses specified in a rich language with hierarchical exception handling, context-dependent methods, etc.

- Built-in constructs that support reflex reactions to fault conditions
- An ability to draw on external expertise for monitoring, diagnosis, and recovery to handle more complex conditions and recovery plans (Pell *et al.* 1998b)
- A high-level control loop that can abort a plan and re-plan, or robustly achieve a context-appropriate standby mode
- A fail-operational capability to restore activities that fail without breaking plans, when possible. This uses recovery methods, and can draw on the property-lock manager to help formulate constraints for use by external recovery planners
- A reusable framework for structuring system-level fault protection in the s/c domain, with pre-defined threads for the major subsystems (power, thermal control, communications, attitude control, etc.)

## General Architectural Properties

This section discusses the general architectural properties necessary for design, deployment, and operation of an executive within a complex mission. Important properties include:

- Modular architecture
- Variable autonomy
- Robustness
- Tools for design, analysis, visualization, communication and logging
- Component testing

## Modular Architecture

One major change underway to EXEC is to modularize it by separating the individual functions, while maintaining and enhancing its functionality. Advantages of a modular architecture include the following:

- The roles and relationships of the components are more cleanly defined. This makes it easier to understand the behavior and interface of each executive capability, and at the implementation level discourages subtle interactions that could otherwise occur.
- The individual components may be chosen in a plug-and-play manner to build up a customized architecture for a given application domain. This reduces the overhead when an application needs only one component, e.g., the resource manager. Plug-and-play also creates an open architecture so that other developers can produce systems satisfying the same capabilities as EXEC components.

- A system that uses executive capabilities no longer needs to be within the framework and language of the executive. When each capability is separated and has a clean interface, a system designer can use the capability with calls, written in any language, to the executive through the interface.
- Once each component capability is better understood, the opportunities for extending and improving the component are more clearly evident. Thus, the overall architecture can be upgraded in a more coherent and more useful fashion.
- Tools for the executive can be specialized to the individual components. This will enhance the effectiveness of each component, as it will allow better formal analysis, visualization, testing, and debugging of each component capability.

Our new modular executive under development, called the Intelligent Deployable Executive Architecture (IDEA), will overcome a major barrier to adoption of executive technologies within conservative space missions.

## Variable Autonomy

An important executive property is to be able to share control of the s/c with ground personnel (Pell *et al.* 1998c). The goal of variable autonomy is to minimize the necessity but maximize the capability of human interaction with a s/c. This increases the confidence of managers and operators responsible for multi-million dollar s/c missions to use an autonomous executive. It also enables operators to deliver new autonomous capability incrementally over the course of a mission. Moreover, operators may desire fully autonomous operation during some phases (like cruising to a target), but little or no autonomous operation during others (like encountering a target). The ability to draw on human expertise, especially in anomalous conditions, can also simplify the design of the system and increase the chance of mission success.

The need to support variable autonomy in addition to high levels of autonomy raises many additional design challenges. Major sources of complexity are that the human operators and the autonomous system have difficulty predicting each other's state and intentions, and their activities may interfere. This is a major area of ongoing research (Pell *et al.* 1998c; Bonasso, Kortenkamp, & Whitney 1997).

## Robustness

In addition to supporting system-level fault protection, an executive must be internally robust in the face of hardware and software faults. Each component must be sensitive to the limits of its own domain of expertise and respond appropriately when these limits are reached. An example is that EXEC tracks the number of times plan execution has failed, and when this count exceeds a threshold EXEC ceases to request new

plans and requests help from ground. Similarly, EXEC checks that the current state of the s/c is within the competence of the planner before requesting a plan.

## Tool Support

For each EXEC component, tools are needed for design, analysis, visualization, communication, and logging. These tools are in various stages of development for different EXEC components. We refer to some of them here.

Simmons & Whelan (1997) has developed a plan execution visualizer (ExecView) compatible with our plan execution system. ExecView enables users to view the plan as it was executed, scroll through the history, and inspect individual activities to get explanations of constraints on these activities. An interesting feature is that the system supports explanations of why an activity was unable to start before a certain time.

Researchers in the Formal Methods group at NASA Ames have worked with us to formally validate some of EXEC's capabilities. Formal analysis of models of the abstract resource management component of EXEC (Gat & Pell 1998) has established the correctness of some aspects of the system, and also located implementation problems earlier in the process (Havelund, Lowry, & Penix 1997; Penix, Alexander, & Havelund 1997; Lowry, Havelund, & Penix 1997). We are planning to extend this work to validate the new resource management component and the plan-runner.

An interesting issue in monitoring and logging behavior of execution systems is that observation overhead can alter the behavior of such time-sensitive systems. We have developed a method for monitoring and logging EXEC *without* significant impact on the runtime behavior of EXEC. The monitoring task is run separately at a relatively low priority. When a high priority task generates logging data, that information is passed in raw form to a "logging" task that will only process the data when time and resources allow. An additional complexity arises in missions with limited communication bandwidth. Such systems often need to support dynamic prioritization of monitoring data.

## Component Testing

Testing an autonomous system is a major challenge because of the wide range of activity it can demonstrate. A problem with systems that carry out behavior over extended time (such as a plan-execution system running a month-long plan) is that testing in real time can be prohibitively expensive. EXEC can scale its clock to run at some fraction of real time and suspend time and resume it at a later date. An enhanced capability, called *warping*, enables EXEC to recognize idle periods and jump directly to the time when a next planned event is supposed to occur.

## Related Work

In this section we discuss work on other execution systems and components.

EXEC is currently implemented on top of ESL (Gat 1996), and a previous version was implemented in RAPS (Firby 1978). These systems both provide languages to support multi-threaded task decomposition, memory databases, and event-waiting and response mechanisms. RAPS provides a tighter level of structuring, which simplifies tracing and potentially generation of action networks from planners. However, ESL provides better support for structured objects and language extension, both of which were crucial for our large software engineering project. PRS (Georgeff & Lansky 1987) and RPL (McDermott 1991) are similar languages that support general execution and agent construction. Interrap (Muller & Pischel 1994) and Golog (Levesque *et al.* 1997) provide execution languages based on logic programming.

Unlike most other general execution systems, TCA (Simmons 1990) explicitly supports an *executive services* perspective. User-level procedures, which can be written in different languages and run on distributed platforms, can access centralized services (such as synchronization, simple resource management, and task decomposition) through well-defined interfaces. While our previous approach was to build integrated execution systems which would explicitly coordinate external software components, we are now moving EXEC to a services perspective as well. EXEC supports a higher-level set of individual capabilities than TCA. Moreover, we are taking the services perspective one step further by providing the ability for each service to be used in a stand-alone fashion and customized according to the needs of an application.

A number of systems have been developed that support specific executive capabilities described in this paper. The attitude and articulation control subsystem (AACS) on the Cassini s/c (Brown, Bernard, & Rasmussen 1995; Hackney, Bernard, & Rasmussen 1993) has explicit software modules for context-dependent command decomposition, resource management, configuration management, and fault protection. While these components are restricted to the particular needs of the mission and the system has no support for closed-loop execution of flexible plans, these Cassini executive capabilities served as a baseline for those incorporated into EXEC.

CIRCA (Musliner, Durfee, & Shin 1993) supports fault protection in the context of guaranteeing mission-critical real-time behavior. CIRCA considers a set of states, actions, and critical failures to be avoided and constructs a program consisting of a set of sense-act transitions which is then executed by a real-time controller. The execution is guaranteed to avoid failure states.

The CONFIG system (Malin 1997; Malin & Leifker 1991), developed at NASA Johnson Space Center,

specifically supports goal-oriented activity definitions and configuration management via declarative models. The Livingstone system (Williams & Nayak 1996) is also designed to function as a model-based manager for discrete hardware configurations. Livingstone's modeling language is based on concurrent transition diagrams. These provide an elegant representation of hardware devices, though the current system does not treat time-delays within repair plans or support procedural event response definitions, as does EXEC.

Boddy (1996) describes a constraint-based distributed scheduling process for air traffic control. Each designated region of airspace is managed by a separate resource manager that allocates spatio-temporal windows to pilots requesting the resource. Musliner & Boddy (1997) describe an application of similar ideas to task distribution for distributed processing. These agents support a similar capability to the modular resource manager we are developing.

Several systems have been developed to support closed-loop plan execution. In contrast with EXEC's current plan execution component, the approach taken in 3T (Bonasso *et al.* 1997) has the planner watch over each step of execution. Hence the planner itself serves as an integral participant in the plan execution capability. Bresina *et al.* (1996) describes APA, which has separate components for generation and execution of temporal plans. However, their approach is currently restricted to single resource domains with no concurrency. Reece & Tate (1994) developed an execution agent for the O-Plan (Currie & Tate 1991) planning system. The combined system supports a plan repair mechanism (Drabble, Tate, & Dalton 1996) that is more sophisticated than that supported by EXEC at present, as it allows the planner to edit any unexecuted portion of the currently executing plan. Our redesigned plan execution component will support a similar editing capability, based on the work in O-Plan and also in Cypress (Wilkins *et al.* 1995). Finally, Lockheed's Tactical Planning and Execution System (TPES) (Mitchell 1997) is an interesting related system that supports many execution and replanning capabilities with a high level of human interaction.

## Summary and Future Work

The Remote Agent Executive demonstrates the capabilities necessary to autonomously control a complex software system. In particular, it supports flexible plan execution, resource management, configuration management, a high-level action-definition language, and system-level fault protection.

In developing EXEC, we have found a number of properties that are desirable for all of the components of an executive. These properties include modular design, variable autonomy, robustness, tools for design, analysis, and visualization, and support for testing.

Our current efforts are focused on decomposing EXEC into separate components, as well as enhancing

each component. The effort to separate out component capabilities is starting with the resource manager, to which is being added the ability to handle resources as continuous quantities and resource reservations. Other component capabilities will be separated out as their roles and relationships become clarified and as applications require.

Our future efforts involve extending the capabilities of the execution agents in a number of directions to support new kinds of missions, increased reliability, and increased resource utilization. These efforts revolve around support for more flexible representations, distributing execution capabilities, coordination among multiple agents, and improvements in development and verification tools.

Other key directions include increased real-time performance and exploiting the structure of plans to gracefully degrade objectives in the face of failures without aborting the current plans.

To conclude, we have described execution capabilities and architectural properties that are important for developing, deploying, and operating autonomous spacecraft. We believe that these capabilities are required for most autonomous agents, and that executives and execution capabilities will become more essential as we attempt to develop autonomous agents of increasing capability.

## Acknowledgments

We gratefully acknowledge the contributions of other members of the DS-1 Remote agent team who influenced the design of EXEC: Abdullah Aljabri, Doug Bernard, Ed Gamble, Erann Gat, Ron Keesing, Sandy Krasner, and Reid Simmons.

## References

- Aljabri, A.; Bernard, D.; Dvorak, D.; Pell, B.; and Starbird, T. 1998. Infusion of autonomy technology into deep space missions: Lessons learned. In *Proceedings of the IEEE Aerospace Conference*. To Appear.
- Bernard, D.; Dorais, G. A.; Fry, C.; Gamble, E. B., Jr.; Kanefsky, B.; Kurien, J.; Millar, W.; Muscettola, N.; Nayak, P. P.; Pell, B.; Rajan, K.; Rouquette, N.; Smith, B.; and Williams, B. C. 1998. Design of the remote agent experiment for spacecraft autonomy. In *Proceedings of the IEEE Aerospace Conference*. To Appear.
- Boddy, M. S. 1996. Contract-based distributed scheduling for a next generation air traffic management system. Technical report, Honeywell Technology Center.
- Bonasso, R. P.; Kortenkamp, D.; Miller, D.; and Slack, M. 1997. Experiences with an architecture for intelligent, reactive agents. *J. Experimental and Theoretical AI* 9(1).
- Bonasso, R. P.; Kortenkamp, D.; and Whitney, T. 1997. Using a robot control architecture to automate space station shuttle operations. In *Procs. of AAAI-97*.
- Bresina, J.; Edgington, W.; Swanson, K.; and Drummond, M. 1996. Operational closed-loop observation scheduling and execution. In Pryor (1996).

- Brown, G.; Bernard, D.; and Rasmussen, R. 1995. Attitude and articulation control for the cassini spacecraft: A fault tolerance overview. In *14th AIAA/IEEE Digital Avionics Systems Conference*.
- Currie, K., and Tate, A. 1991. O-Plan: the open planning architecture. *Art. Int.* 52(1):49–86.
- Drabble, B.; Tate, A.; and Dalton, J. 1996. O-plan project evaluation experiments and results. O-Plan Technical Report ARPA-RL/O-Plan/TR/23 Version 1, AIAI.
- Firby, R. J. 1978. *Adaptive execution in complex dynamic worlds*. Ph.D. Dissertation, Yale University.
- Gat, E., and Pell, B. 1998. Abstract resource management in an unconstrained plan execution system. In *Proceedings of the IEEE Aerospace Conference*. To Appear.
- Gat, E. 1996. ESL: A language for supporting robust plan execution in embedded autonomous agents. In Pryor (1996).
- Georgeff, M. P., and Lansky, A. L. 1987. Procedural knowledge. Technical Report 411, Artificial Intelligence Center, SRI International.
- Hackney, J.; Bernard, D.; and Rasmussen, R. 1993. The Cassini spacecraft: Object oriented flight control software. In *16th Annual AAS Guidance and Control Conference*.
- Havelund, K.; Lowry, M.; and Penix, J. 1997. Formal analysis of a spacecraft controller using SPIN. Technical report, NASA Ames Research Center. In preparation.
- Levesque, H.; Reiter, R.; Lesperance, Y.; Lin, F.; and Scherl, R. 1997. Golog: A logic programming language for dynamic domains. *Journal of Logic Programming* 31:59–84.
- Lowry, M.; Havelund, K.; and Penix, J. 1997. Verification and validation of ai systems that control deep-space spacecraft. In *Proc. ISMIS'97*.
- Malin, J. T., and Leifker, D. B. 1991. Functional modeling with goal-oriented activities for analysis of effects of failures on functions and operations. *Informatics & Telematics* 8(4):353–364.
- Malin, J. T. 1997. Statement of interest: Designing model-based autonomous systems for coordinated acquisition and maintenance of models and procedures. In Nayak and Williams (1997).
- McDermott, D. 1991. A reactive plan language. Technical report, Computer Science Dept, Yale University.
- Mitchell, S. W. 1997. A hybrid architecture for real-time mixed-initiative planning and control. In *Procs. of AAAI-97*, 1032–1037.
- Muller, J., and Pischel, M. 1994. An architecture for dynamically interacting agents. *International Journal of Intelligent and Cooperative Information Systems (IJICIS)* 3(1):25–45.
- Muscettola, N.; Pell, B.; Hansson, O.; and Mohan, S. 1995. Automating mission scheduling for space-based observatories. In Henry, G., and Eaton, J., eds., *Robotic Telescopes: Current Capabilities, Present Developments, and Future Prospects for Automated Astronomy*, number 79 in ASP Conf. Series. Provo, UT: Astronomical Society of the Pacific.
- Muscettola, N.; Morris, P.; Pell, B.; and Smith, B. 1998. Issues in temporal reasoning for autonomous control systems. In Wooldridge (1998). To appear.
- Musliner, D., and Boddy, M. S. 1997. Contract-based distributed scheduling for distributed processing. In *Working Notes of the AAAI Workshop on Constraints and Agents*.
- Musliner, D.; Durfee, E.; and Shin, K. 1993. Circa: A cooperative, intelligent, real-time control architecture. *IEEE Transactions on Systems, Man, and Cybernetics* 23(6).
- Nayak, P. P., and Williams, B. C., eds. 1997. *Proceedings of the AAAI Fall Symposium on Model-Directed Autonomous Systems*. AAAI Press.
- Pell, B.; Gat, E.; Keesing, R.; Muscettola, N.; and Smith, B. 1997. Robust periodic planning and execution for autonomous spacecraft. In *Procs. of IJCAI-97*.
- Pell, B.; Bernard, D. E.; Chien, S. A.; Gat, E.; Muscettola, N.; Nayak, P. P.; Wagner, M. D.; and Williams, B. C. 1998a. An autonomous spacecraft agent prototype. *Autonomous Robotics* 5(1). To Appear.
- Pell, B.; Gamble, E.; Gat, E.; Keesing, R.; Kurien, J.; Millar, B.; Nayak, P. P.; Plaunt, C.; and Williams, B. 1998b. A hybrid procedural/deductive executive for autonomous spacecraft. In Wooldridge (1998). To appear.
- Pell, B.; Sawyer, S.; Bernard, D. E.; Muscettola, N.; and Smith, B. 1998c. Mission operations with an autonomous agent. In *Proceedings of the IEEE Aerospace Conference*. To Appear.
- Penix, J.; Alexander, P.; and Havelund, K. 1997. Declarative specification of software architectures. In Lowry, M., ed., *Proceedings of the International Conference on Automated Software Engineering*. IEEE Computer Society Press.
- Pryor, L., ed. 1996. *Proceedings of the AAAI Fall Symposium on Plan Execution*. AAAI Press.
- Reece, G., and Tate, A. 1994. Synthesizing protection monitors from causal structure. In *Procs. AIPS-94*. AAAI Press.
- Simmons, R., and Whelan, G. 1997. Visualization tools for validating software of autonomous spacecraft. In Atkinson, D., ed., *Proceedings of the Fourth International Symposium on Artificial Intelligence, Robotics, and Automation for Space (i-SAIRAS)*. Tokyo, Japan: Jet Propulsion Laboratory.
- Simmons, R. 1990. An architecture for coordinating planning, sensing, and action. In *Procs. DARPA Workshop on Innovative Approaches to Planning, Scheduling and Control*, 292–297. San Mateo, CA: DARPA.
- Wilkins, D. E.; Myers, K. L.; Lowrance, J. D.; and Wesley, L. P. 1995. Planning and reacting in uncertain and dynamic environments. *Journal of Experimental and Theoretical AI* 7(1):197–227.
- Williams, B. C., and Nayak, P. P. 1996. A model-based approach to reactive self-configuring systems. In *Procs. of AAAI-96*, 971–978. Cambridge, Mass.: AAAI.
- Wooldridge, M., ed. 1998. *Proceedings of the Second International Conference on Autonomous Agents*. ACM Press. To appear.