Experiments with Diversified Models for Fault-Tolerant Planning

Benjamin Lussier, Matthieu Gallien, Jérémie Guiochet,
Félix Ingrand, Marc-Olivier Killijian, David Powell
LAAS-CNRS, 7 avenue du Colonel Roche, 31077 Toulouse Cedex 4, France
firstname.lastname@laas.fr

***Abstract*—Autonomous robots make extensive use of decisional mechanisms, such as planning. These mechanisms are able to take complex and adaptative decisions, but are notoriously hard to validate. This paper reports an investigation of how redundant, diversified models can be used to tolerate residual design faults in such mechanisms. A fault-tolerant temporal planner has been designed and implemented using diversity, and its effectiveness demonstrated experimentally through fault injection. The paper describes the implementation of the fault-tolerant planner and discusses the results obtained. The results indicate that diversification provides a noticeable improvement in planning reliability with a negligible performance overhead. However, further improvements in reliability will require implementation of a on-line checking mechanism for assessing plan validity before execution.**

I. Introduction

Complex autonomous systems generally rely on planning for selecting and organizing actions to achieve high-level goals fixed by the user. Experimental systems using planning as a decisional mechanism have already been implemented as space exploration satellites, space rovers, museum tour guides, and personal assistants. However, such systems are not yet ready for real life utilization, as the dependability of their decisional mechanisms remains in question. Indeed, how can we justifiably trust planning mechanisms, whose behavior is difficult to predict and validate? To increase the confidence that we may have in such mechanisms so that they may be used in critical applications, we propose in this paper a fault tolerance approach focused on development faults in the planner knowledge. The approach uses redundant diversified planning models.

First, we present fault tolerance techniques that may be applied to planning mechanisms and the implementation of a fault tolerant planner for an autonomous system architecture. Second, we introduce the validation framework that we developed to assess performance and efficacy of our fault tolerance approach. Finally, we present the experimental results produced during our validation.

II. Fault Tolerant Planning

We propose here fault tolerant techniques adapted to planners and focusing on development faults in the planner knowledge. We first present issues faced by dependability techniques when confronted to planning mechanisms. We then propose four error detection mechanisms and two recovery mechanisms adapted to planners and a redundancy management component that coordinates these mechanisms. We finally present an example of implementation of the proposed component in an existing autonomous system architecture.

A. Dependability Issues

Planning, like other decisional mechanisms, poses significant challenges for validation. Classic problems faced by testing and verification are exacerbated. First, *execution contexts* in autonomous systems are neither controllable nor completely known; even worse, consequences of the system actions are often uncertain. Second, planning mechanisms have to be *validated in the complete architecture*, as they aim to enhance functionalities of the lower levels through high level abstractions and actions. Integrated tests are thus necessary very early in the development cycle. Third, the *oracle problem*[1] is particularly difficult since (a) equally correct plans may be completely different and (b) an unforeseen adverse environmental situation may completely prevent some goals from being achieved, thus ineluctably degrading the system performance, however well it behaves (for example, cliffs, or some other feature of the local terrain, may make a position goal unreachable).

One way to address the latter issue is to define an oracle as a set of constraints that necessarily and sufficiently characterizes a correct plan: plans satisfying the constraints are deemed correct. Such a technique was used for thorough testing of the RAX planner during the NASA *Deep Space One* project [3], or in the VAL validation tool [9]. Extensive collaboration of application and planner experts is necessary to generate the correct set of constraints.

Automatic static analysis may also be used to ascertain properties on planning models, whereas manual static analysis requires domain experts to closely scrutinize models proposed by planning developers. For example, the development tool Planware [2] offers facilities for both types of analysis. A Failure Recovery Analysis tool is proposed in [8] to ease model corrections during development.

Some work has also been done on evaluating planning dependability. A measure for planner reliability is proposed in [5], which compares theoretical results to experimental ones,

[1] How to conclude on correctness of a program's outputs to selected test inputs?

showing a necessary compromise between temporal failures (related to calculability of decisional mechanisms) and value failures (related to correctness of decisional mechanisms). Later work [4] proposes concurrent use of planners with diversified heuristics to answer this compromise: a first heuristic, quick but dirty, is used when a slower but more focussed heuristic fails to deliver a plan in time. To our knowledge, no other fault tolerance mechanisms have been proposed in this domain. We strongly believe, however, that such mechanisms are essential to provide more dependability in autonomous systems.

B. Principles

Complementary to testing, diversity is the only known approach to improve trust in the behavior of a critical system regarding development faults (e.g., diversification is used in software components of the Airbus A320, and in hardware components of the Boeing B777). The general principle of the mechanisms that we propose is similar to that of recovery blocks [16] and distributed recovery blocks [11]. Diversified variants of the planner are executed either sequentially or concurrently. A recovery choice is made according to errors detected in the plans produced. Diversity between the variants is encouraged by forcing the use of different algorithms, variable domains and parameters in the models, and different heuristics of the variants.

1) Detection: Implementing error detection for decisional mechanisms in general, and planners in particular, is difficult [14]. There are often many different valid plans, which can be quite dissimilar. Therefore, error detection by comparison of redundantly-produced plans is not a viable option. Thus, we must implement error detection by independent means. Here, we propose four complementary error detection mechanisms: a *watchdog timer*, a *plan analyzer*, a *plan failure detector* and an *on-line goal checker*.

A watchdog timer is used to detect when the search process is too slow or when a critical failure such as a deadlock occurs. Timing errors can be due to faults in the planner model, in its search engine, or ineffectiveness of the search heuristics.

A plan analyzer can be applied on the output of the planner. It is an acceptance test (i.e., an on-line oracle) that verifies that the produced plan satisfies a number of constraints and properties. This set of constraints and properties can be obtained from the system specification and from domain expertise but it must be diverse from the planner model. This mechanism is able to detect errors due to faults in the planner model or heuristics, and in the search engine.

A plan failure detector is a classical mechanism used in robotics for execution control. Failure of an action that is part of the plan may be due to an unresolvable adverse environmental situation, or may indicate errors in the plan due to faults in the knowledge or in the search engine. Usually, when such an action failure is raised, the search engine tries to repair the plan. When this is not possible, it raises a plan failure. We use these plan failure reports for detection purposes.

```
1. begin mission
2.   failed_planners ← ∅;
3.   while (goals ≠ ∅)
4.     candidates ← planners;
5.     while (candidates ≠ ∅ & goals ≠ ∅)
6.       choose k such as (k ∈ candidates)
                         & (k ∉ failed_planners);
7.       candidates ← candidates \ k;
8.       init_watchdog(max_duration);
9.       send (plan_request) to k;
10.      wait % for either of these two events
11.        □ receive (plan) from k
12.          stop_watchdog;
13.          if analyze(plan)=OK then
14.             failed_planners ← ∅;
15.             k.execute_plan();
                % if the plan fails goals != empty
                % and then we loop to line 5 or 3
16.          else
17.             send(k.invalid_plan) to operator;
18.             failed_planners ← failed_planners ∪ k;
19.          end if
20.        □ watchdog_timeout
21.          failed_planners ← failed_planners ∪ k;
22.      end wait
23.      if failed_planners = planners then
24.        raise exception "no valid plan
                          found in time";
           % no remaining planner,
           % the mission has failed
25.      end if
26.    end while
27.  end while
28. end mission
```

Fig. 1. Sequential Planning Policy

An on-line goal checker verifies whether goals are reached while the plan is executed. Goals can only be declared as failed when every action of the plan has been carried out. This implies that the checker maintains an internal representation of the system state and of the goals that have been reached.

2) Recovery: We propose two recovery mechanisms, both using different planners based on diverse knowledge. Theoretically, different planners could be used to tolerate not only faults in planning models but also in the search engine. However, no practical work has been done to test this possibility.

With the first mechanism, the planners are executed sequentially, one after another. The principle is given in Figure 1. Basically, each time an error is detected, we switch to another planner until all goals have been reached or until all planners fail in a row. Once all the planners have been used and there are still some unsatisfied goals, we go back to the initial set of planners. This algorithm illustrates the use of the four detection mechanisms presented in Section II-B1: watchdog timer (lines 8 and 20), plan analyzer (line 13), plan failure detector (line 15), on-line goal checker (lines 3 and 5).

Reusing planners that have been previously detected as failed makes sense for two different reasons: (a) a perfectly correct plan can fail during execution due to an adverse environmental situation, and (b) some planners, even faulty, can still be efficient for some settings since the situation that activated the fault may have disappeared.

It is worth noting that the choice of the planners, and the order in which they are used, is arbitrary in this particular example (line 6). However, the choice of the planner could

```
1.  begin mission
2.    while (goals ≠ ∅)
3.      candidates ← planners;
4.      init_watchdog(max_duration);
5.      send (plan_request) to candidates;
6.      while (candidates ≠ ∅)
7.        wait % for either of these two events
8.        □ receive (plan) from k ∈ candidates
9.          candidates ← candidates \ k;
10.         pause_watchdog;
11.         if analyze(plan)=OK then
12.           stop_watchdog;
13.           send (cancel_planning) to candidates;
14.           candidates ← ∅;
15.           k.execute_plan();
              % if the plan fails goals != empty
              % and then we loop to line 3
16.         else
17.           resume_watchdog;
18.           send(k.invalid_plan) to operator;
19.           if (candidates = ∅)
20.              raise exception "no valid plan
                                  found in time";
                 % no remaining planner,
                 % the mission has failed
21.            end if
22.          end if
23.       □ watchdog_timeout
24.          raise exception "no valid plan
                              found in time";
             % no remaining planner,
             % the mission has failed
25.       end wait
26.     end while
27.   end while
28. end mission
```

Fig. 2. Concurrent Planning Policy

take advantage of application-specific knowledge about the most appropriate planner for the current situation or knowledge about recently observed failure rates of the planners.

With the second recovery mechanism, presented in Figure 2, the planners are executed concurrently [13]. The main differences with respect to the algorithm given in Figure 1 are: (a) the plan request message is sent to every planning candidate (line 5), (b) when a correct plan is found, the other planners are requested to stop planning (line 13), and (c) a watchdog timeout means that all the planners have failed (line 23).

Here, the choice of planner order is implicit: the first planner obtaining a plan is chosen. However, this could lead to the repeated selection of the same faulty but rapid planner. Some additional mechanism is thus required to circumvent this problem. For example, the planner selected during the previous round can be withdrawn from the set of candidates for the current round.

3) Coordination: From a dependability point of view, the fault-tolerance mechanisms have to be as independent as possible from the decisional layer, i.e., in this case from the planners. This is why we propose to handle both the detection and recovery mechanisms, and the services necessary for their implementation, in a middleware level component called FTplan, standing for *Fault-Tolerant PLANner coordinator*.

This component has to integrate the fault tolerance mechanisms into the robot architecture. This implies essentially communication between, and synchronization and coordination of, the error detection mechanisms and the redundant planners.

To avoid error propagation from a possibly faulty planner, FTplan should not take any information that comes from or depends on the planners themselves. The watchdog can easily be implemented from the operating system timing primitives. Action and plan failure detection are performed at the lower plan execution control layer, so error reports can be obtained and reused by FTplan. A plan analyzer performs simple acceptance checks using rules expressed independently from the planners and their knowledge.

However, implementing an on-line goal checker without relying on information obtained through the planner is more difficult. FTplan maintains for this purpose its own system state representation, based on information obtained from the plan execution control layer. This system state representation is checked against the set of goals prescribed for the current mission.

Whatever the particular recovery mechanism it implements, sequential or parallel, FTplan has to manage several planners. It needs to communicate with them, e.g., for sending plan requests or for updating their goals and system state representations before replanning. It also needs to be able to control their life cycle: start a new instance or even stop one when it takes too long to produce a plan.

FTplan is intended to allow tolerance of development faults in planners (and particularly in planning models). FTplan itself is *not* fault-tolerant, but being much simpler than the planners it coordinates, we can safely rely on classic verification and testing to assume that it is fault-free.

C. Implementation

We present here the implementation of the proposed mechanisms. We introduce the target architecture and then give some implementation details about the FTplan component.

1) LAAS Architecture: The LAAS architecture is presented in [1], and some recent modifications have been proposed in [12]. It has been successfully applied to several mobile robots, some of which have performed missions in real situations (human interaction or exploration). It is composed of three main components[2] as presented in Figure 3: GenoM modules, OpenPRS, and IxTeT.

The functional level is composed of a set of automatically generated *GenoM modules*, each of them offering a set of services, which perform computation (e.g., trajectory movement calculation) or communication with physical devices (sensors and actuators).

The procedural executive *OpenPRS* (*Open Procedural Reasoning System*), is in charge of decomposing and refining plan actions into lower-level actions executable by functional components, and executing them. This component links the decisional component (IxTeT) and the functional level. During execution, OpenPRS reports any action failures to the planner, in order to re-plan or repair the plan. As several IxTeT

[2] An additional robustness component, R2C, is introduced in [15]. We have not considered it in this study since its current implementation is not compatible with our experimentation environment.

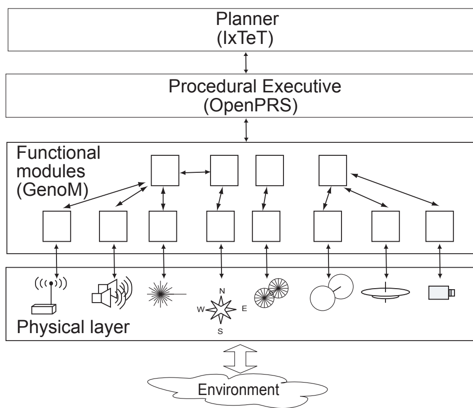

Fig. 3. The LAAS architecture

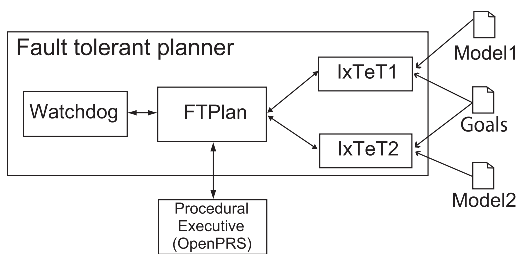

Fig. 4. Fault tolerant planner

actions can be performed concurrently, it has also to schedule sequences of refined actions.

IxTeT (*IndeXed TimE Table*) is a temporal constraint planner, combining high level actions to build plans. Each *action* is described in a *model* file used by the planner as a set of constraints on attributes (e.g., robot position), resources (e.g., energy consumption), numeric or temporal data (e.g., action duration). Then, a valid plan is calculated combining a set of actions in such a way that they are conflict-free and they fulfill the goals. The description of actions in the planner model is critical for the generation of successful plans and thus for the dependability of the robot as a whole.

2) Fault Tolerant Planner Implementation: The fault tolerance principles presented in Section II-B have been implemented in a fault tolerant planner component as presented in Figure 4. This component replaces the original component "Planner" presented in Figure 3. The FTplan component is in charge of communicating with OpenPRS as the original planner does.

The current version of FTplan implements the sequential redundant planner coordination algorithm presented earlier (Section II-B, Figure 1) with two IxTeT planners. Currently, the plan analysis function is empty (it always return *true*) so error detection relies solely on just three of the mechanisms presented in Section II-B1: watchdog timer, plan failure detection, and on-line goal checker.

The watchdog timer is launched at each start of planning. As soon as a plan is found before the assume worst-case time limit (40 seconds in our implementation), the watchdog is stopped. If timeout occurs, FTplan stops the current IxTeT, and sends a plan request to the other IxTeT planner, until a plan is found or both planners have failed. In the latter case, the system is put in a safe state (i.e., all activities are ceased), and an error message is sent to the operator.

On-line goal checking is performed after each action executed by OpenPRS that can result in a modification in the goal achievements (for instance, in the application considered in Section III-B: a camera shot, a communication, movement of the robot, etc.). This checking is carried out by analyzing the system state at the end of an action, determining goals that may have been accomplished and checking that no inconsistent actions have been executed simultaneously. Unfulfilled goals are resubmitted to the planner during the next replanning or at the end of plan execution.

In the current implementation, FTplan checks every 10ms if there is a message from OpenPRS or one of the IxTeT planners. In case of an action request from a planner or an action report from OpenPRS, FTplan updates its system representation before transmitting the request. If the request is a plan execution failure (the system has not been able to perform the actions of the plan), then FTplan launches a replan using the sequential mechanism. If the request indicates that the actions are finished, then FTplan checks if the goals have been reached.

III. Experiments and Validation

Our validation framework relies on simulation and fault injection. Simulation is used since it is both safer and more practical to exercise the autonomy software on a simulated robot than on a real one. Fault injection is used since it is the *only* way to test the fault tolerance mechanisms with respect to their specific inputs, i.e., faults in planning knowledge. In the absence of any evidence regarding real faults, there is no other practical choice than to rely on *mutations*[3], which have been found to efficiently simulate real faults in imperative languages [7].

We now introduce successively the targeted software architecture, the workload, the faultload, and the readouts and measurements we obtain from system activity.

A. Software Architecture

Our simulation environment is represented in Figure 5. It incorporates three elements: an open source robot simulator named Gazebo, an interface library named Pocosim, and the components of the LAAS architecture already presented in section II-C1.

The *robot simulator Gazebo*[4] is used to simulate the physical world and the actions of the autonomous system; it takes as input a file describing the environment of the simulation (mainly a list of static or dynamic obstacles containing their

[3] A mutation is a syntactic modification of an existing program.
[4] "The player/stage project", http://playerstage.sourceforge.net

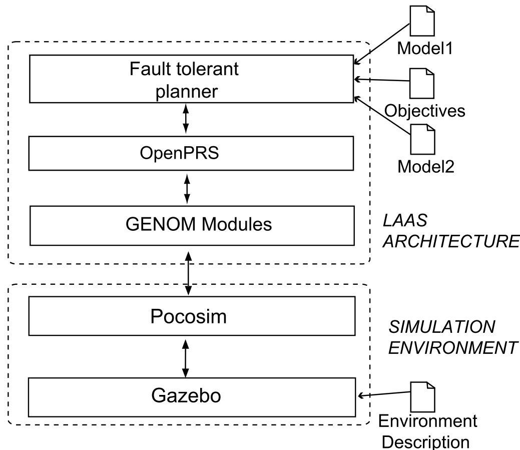

Fig. 5. Simulation environment

position, and the physical description of the robot) and executes the movement of the robot and dynamic obstacles, and possible interactions between objects.

The *Pocosim library* [10] is a software bridge between the simulated robot (executed on Gazebo) and the software commands generated by the GenoM modules: it transforms commands to the actuators into movements or actions to be executed on the simulated robot, and relays the sensor inputs that Gazebo produces from the simulation.

Our *autonomous system* is based on an existing ATRV (All Terrain Robotic Vehicle) robot, and employs GenoM software modules interfaced with the Gazebo simulated hardware. The upper layer of the LAAS architecture executes as presented in the previous section. Two different models are used with the IxTeT planners. The first model was thoroughly tested and used on a real ATRV robot; we use it as primary model and as target for fault injection. We specifically developed the second model through forced diversification of the first: for example, the robot position is characterized numerically in the first model and symbolically in the second.

B. Workload

Our workload mimics the possible activity of a space rover. The system is required to achieve three subsets of goals: *take science photos* at specific locations (in any order), *communicate* with an orbiter during specified visibility windows, and *be back at the initial position* at the end of the mission.

To partially address the fact that the robot must operate in an open unknown environment, we chose to activate the system's functionalities in some representative situations resulting from combinations of sets of missions and worlds. A *mission* encompasses the number and location of photos to be taken, and the number and occurrence of visibility windows. A *world* is a set of static obstacles unknown to the robot (possibly blocking the system from executing one of its goals), which introduces uncertainties and stresses the system navigation mechanism.

We implemented four missions and four worlds, thus applying sixteen execution contexts to each mutation. Missions are referenced as gradually more difficult M1, M2, M3 and M4: M1 consists in two communications and three photos in close locations, whereas M4 consists in four communications and five far apart photos. Environments are referenced as worlds W1, W2, W3 and W4. W1 is an empty world, with no obstacles to hinder plan execution. W2 and W3 contains small cylindrical obstacles, whereas W4 includes large rectangular obstacles that may pose great difficulties to the navigation module, and are susceptible to endlessly block the robot path.

In addition, several equivalent experiments are needed to address the non-determinacy of the experiments. This is due to asynchrony in the various subsystems of the robot and in the underlying operating systems: task scheduling differences between similar experiments may degrade into task failures and possibly unsatisfied goals, even in the absence of faults. We thus execute each basic experiment three times, leading to a total of 48 experiments per mutation. More repetition would of course be needed for statistical inference on the basic experiments but this would have led to a total number of experiments higher than that which could have been carried out with the ressources available (each basic experiment lasts about 20 minutes).

C. Faultload

To assess performance and efficacy of the proposed fault tolerance mechanisms, we inject faults in a planning model by random mutation of the model source code (i.e., in Model1 of Figure 5). Five types of possible mutations were identified from the model syntax:

1) Substitution of numerical values: each numerical value is exchanged with members of a set of real numbers that encompasses (a) all numerical variables in all the tasks of the model, (b) a set of specific values (such as 0, 1 or -1), and (c) a set of randomly-selected values.
2) Substitution of variables: since the scope of a variable is limited to the task where it is defined, numerical (resp. temporal) variables are exchanged with all numerical (resp. temporal) variables of the same task.
3) Substitution of attribute values: in the IxTeT formalism, attributes are the different variables that together describe the system state. Attribute values in the model are exchanged with other possible values in the range of the attribute.
4) Substitution of language operators: in addition to classic numerical operators on temporal and numerical values, the IxTeT formalism employs specific operators, such as "nonPreemptive" (that indicates that a task cannot be interrupted by the executive).
5) Removal of a constraint relation: a randomly selected constraint on attributes or variables is removed from the model.

Substitution mutations were automatically generated using the SESAME tool [6]. Using an off-line compilation, this tool detects and eliminates binary equivalent or syntactically

incorrect mutants. Removal of random constraint relations was carried out through a PERL script and added to the mutations generated by SESAME. All in all, more than 1000 mutants were generated from the first model.

For better representativeness of injected faults, we consider only mutants that are able to find a plan in at least one mission (we consider that models that systematically fail would easily be detected during the development phase). As a simple optimization, given our limited resources, we also chose to carry out a simple manual analysis aimed at eliminating mutants that evidently could not respect the above criterion.

D. Records and Measurements

Numerous log files are generated by a single experiment: simulated data from Gazebo (including robot position and hardware module activity), output messages from GenoM modules and OpenPRS, requests and reports sent and received by each planner, as well as outputs of the planning process.

Problems arise however in trying to condense this amount of data into significant relevant measures. Contrary to more classic mutation experiments, the result of an experiment cannot be easily dichomotized as either failed or successful. As previously mentioned, an autonomous system is confronted with partially unknown environments and situations, and some of its goals may be difficult or even impossible to achieve in some contexts. Thus, assessment of the results of a mission must be graded into more than just two levels. Moreover, detection of equivalent mutants is complexified by the non-deterministic context of autonomous systems

To answer these issues to some extent, we chose to categorize the quality of the result of an experiment with: (a) the subset of goals that have been successfully achieved, and (b) performance results such as the mission execution time and the distance covered by the robot to achieve its goals. Due to space constraints, we focus in the rest of this paper on measurements relative to the mission goals.

IV. Results

We present in this part several experimental results using the evaluation framework previously introduced. Experiments were executed on i386 systems with a 3.2 GHz CPU and Linux OS. We first study the performance cost of the proposed mechanisms, then present their efficacy in tolerating injected faults through three mutation examples, and then global results for 28 injected faults.

A. Fault-free Performance

To determine the overhead of the proposed fault tolerance mechanisms, we first concentrate on supposed fault-free models. Figure 6 presents the impact of FTplan on the system behavior.

Note that results in W4 must be treated with caution, as this world contains large obstacles that may cause navigational failures and block the robot path forever. As our work focuses on planning model faults rather than limitations of functional modules, we consider that success in this world relies more on serendipity in the choice of plan rather than correctness of the planner model. It is however interesting to study the system reaction to unforeseen and unforgiving situations that possibly arise in an open and uncontrolled environment. Note that these results show that different models give rise to different failure behaviors: particularly in W4, the three systems fail differently.

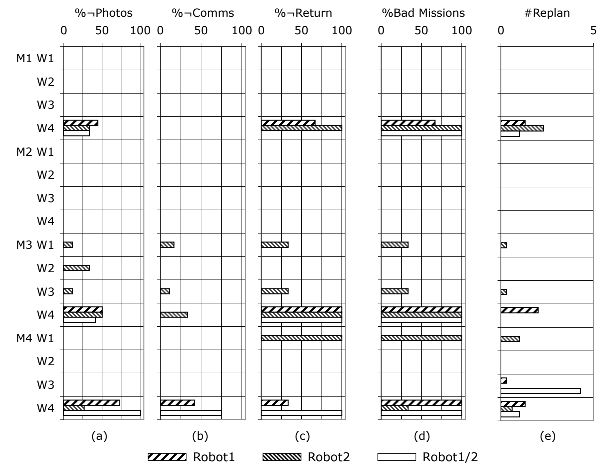

Fig. 6. Impact of FTplan (without injected faults): *This figure studies the impact of the FTplan component on fault-free system behavior by comparing three different robots: Robot1 uses our first model, Robot2 uses our second model, and Robot1/2 contains an FTplan component that uses successively our first and second models. For each considered activity (M1W1 to M4W4), the figure shows five different measures: (a) (b) (c) three failure proportions to reach the different types of goals in a mission (resp. photos, communications, and returns to initial position), (d) failure proportion of the whole mission (a mission is considered failed if one or more mission goals were not achieved), and (e) the mean number of replanning operations observed during one experiment (in the case of Robot1/2, this number is equivalent to the number of model switches during the mission).*

W4 set aside, results are globally very good: Robot1 and Robot1/2 succeed in all their goals, while Robot2 fails a few goals in M3, and all its return goals in M4W1. These failures may be attributed to a larger set of constraints in this model that may be costly in performance, and underestimated distance declarations. The mean activity time of the systems (that is the time until the system stops all activity in a mission) is an average of 404 seconds for Robot1, 376 seconds for Robot2, and 405 seconds for Robot1/2. Time performance-wise, the three systems are thus roughly equivalent.

Although the results are mostly positive, showing that FTplan's main execution loop does not severely decrease goal achievement or performance in the chosen scenarios, they are still insufficient to assess the overhead of planner switches (i.e., replanning requests in our implementation) as very few occurred in these fault-free experiments. However, further experiments involving numerous replanning requests showed that FTplan did not severely degrade the system performance. Due to lack of space, we do not develop this aspect here.

B. Fault-tolerance efficacy

To test the efficacy of the proposed mechanisms and the FTplan component, we injected 38 faults in our first model,

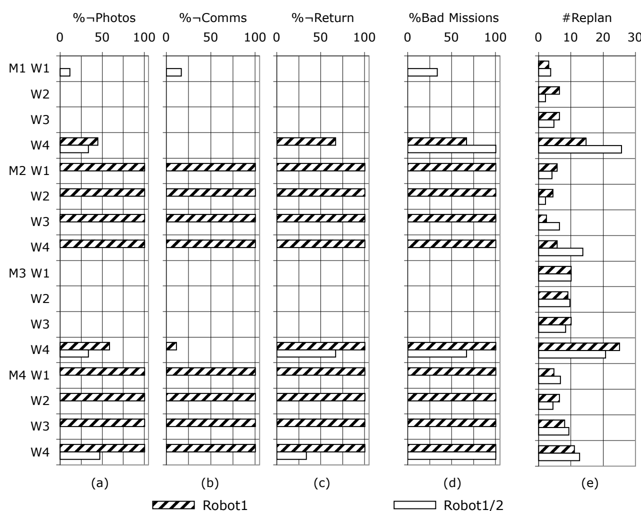

Fig. 7. Results for mutation 1-39

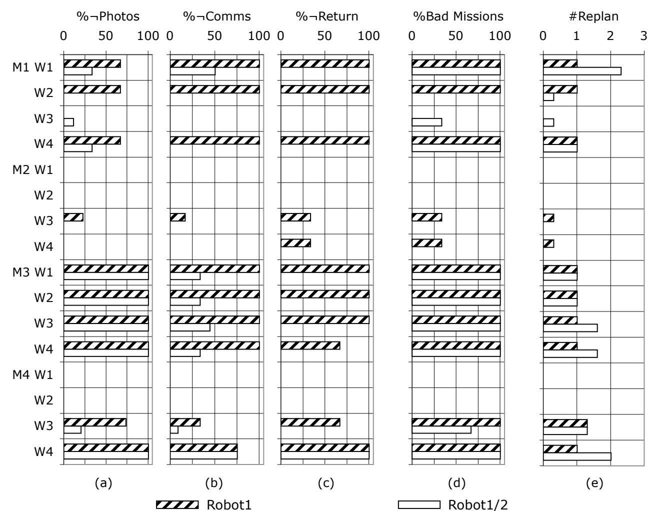

Fig. 8. Results for mutation 1-589

realizing more than 3500 experiments equivalent to 1200 hours of testing. We discarded 10 mutants that were unable to find a plan for any of the four missions[5]. We believe that five of the remaining mutants are equivalent to the fault-free model. However, the non-deterministic nature of autonomous systems makes it delicate to define objective equivalence criteria. We thus include the results obtained with these five mutants, leading to a pessimistic estimation of the improvement offered by FTplan.

The 28 considered mutations are categorized in the following manner: three substitutions of attribute values, six substitutions of variables, ten substitutions of numerical values, four substitutions of operators, and six removals of constraints. The mutants were executed on Robot1 and Robot1/2.

1) Examples of experimental results: We develop here the results of three mutations as examples of our experiments.

Results for the first mutation, identified by the indice 1-39, are presented in Figure 7. This mutation causes an overly constrained robot position during a camera shot. It thus results in the planner's inability to find plans for missions where photographs have to be taken in positions that do not respect the overly-restrictive constraint (this is the case for missions M2 and M4). This example illustrates the significance of the system activity chosen for the evaluation: numerous different missions are necessary because faults can remain dormant in some missions. Since testing the practically infinite execution context is well nigh impossible, this example underlines the difficulty of testing and thus the interest of fault tolerance mechanisms to cover residual faults in the deployed planning models.

Figure 8 presents the results for mutation 1-589. The fault injected in this mutation affects only a movement recovery action of the planning model. Thus, contrary to the previous example, correct plans are established for all missions. However, as soon as a movement action fails, the planner is unable to find a plan allowing recovery of the movement, which causes failure of the system. This is particularly obvious in the case of missions M1 and M3, where short distances between photograph locations lead to a short temporal margin for the action movement. As acceleration, deceleration and rotation are not considered in the planning model, movements are susceptible to take longer than estimated, and can thus be be interrupted by the plan execution controler and considered failed, necessiting a recovery action. In missions M2 and M4, movements cover greater distances, resulting in larger temporal margins and thus fewer movement action failures. Robot1/2 tolerates this fault to some extent: completely in mission M2 and partially in mission M4. The high failure rate of mission M3 for Robot1/2 can be explained by a domino effect due to communication goals being given priority over photography goals. When the fault is activated due to a failed movement action, FTplan switches to Model2 and requests a plan. However, a communication goal is now so near that the planner is unable to find a plan to achieve it, so it abandons goals of lesser priority, but to no avail. This example raises two important comments:

- First, testing with numerous diversified missions and environments is once again pointed out, as the fault is not activated in several activities.
- Second, testing must be realized in an *integrated system*. Indeed, the original plans produced by the planner are correct, as well as the lower levels of the system. However, the planning model contains a serious fault that can cause critical failure of the system in some executions.

The results of mutation 1-583 are presented in Figure 9. It is the only case in our 28 experiments where the fault intolerant Robot1 shows *better* results than the fault tolerant Robot1/2, although we identified 8 other mutations where results in both systems were similar (including the five mutants suspected to be equivalent to the original model). The injected fault causes

[5]In this case, Robot1/2 gives the same results as the fault-free Model2: nearly perfect success rates in W1, W2 and W3.

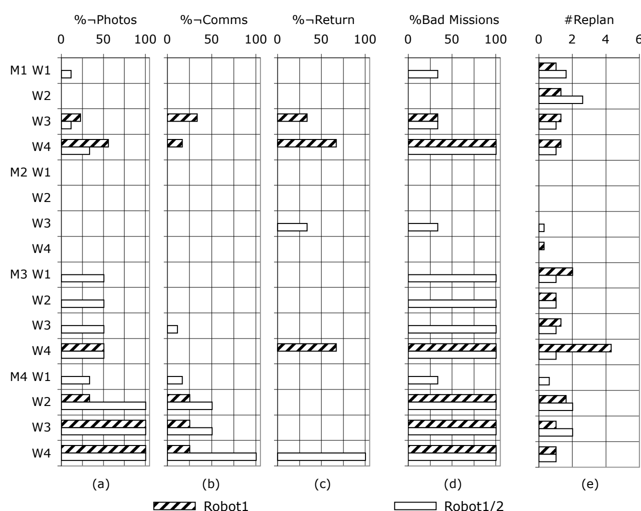

Fig. 9. Results for mutation 1-583

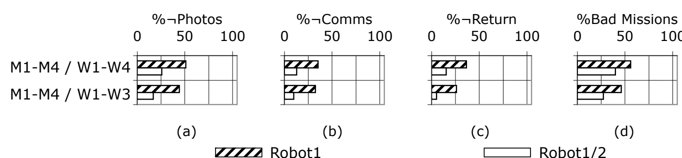

Fig. 10. Impact of planner redundancy (with injected faults): *This figure presents overall results achieved for all 28 mutations, both with and without the heavily-constrained world W4.*

the duration of the robot movement to be underestimated in plans, resulting in execution errors and replannings that lead to some goals being missed. Failures of Robot1/2 are not directly linked to the injected fault, but rather to the use of the badly optimized second model. It uses more pessimistic time constraints than the first model, thus giving up some goals, and causing failure of all photographs in mission M4 through a similar domino effect as that presented in the previous example.

2) General Results: The general results, including all 28 mutations, are presented in Figure 10. These results give objective evidence that model diversification favorably contributes to fault tolerance of an autonomous system considering the proposed faultload: failure decreases for photo goals of 62% (respectively, 50% including W4), 70% (64%) for communication goals, 80% (58%) for returns goals, and 41% (29%) for whole missions. Note, however, that RobotFT in the presence of injected faults is *less* successful than a single fault-free model (cf. Figure 6). This apparent decrease in dependability is explained by the fact that, in our current implementation, incorrect plans are only detected when their execution has at least partially failed, possibly rendering one or more goals unachievable, even after recovery. This underlines the importance of plan analysis procedures to attempt to detect errors in plans *before* they are executed.

V. Conclusion

The work presented in this paper proposes fault tolerant mechanisms based on diversified planning models. We developed a component providing error detection and recovery appropriate for fault-tolerant planning, and implemented it in the LAAS architecture. This component can use four detection mechanisms (watchdog timer, plan failure detector, on-line goal checker and plan analyzer), and two recovery policies (sequential planning and concurrent planning). Our current implementation is that of sequential planning associated with the first three error detection mechanisms.

To assess the performance overhead and the efficacy of the proposed mechanisms, we developed a validation framework that exercises the software on a simulated robot platform, and carried out what we believe to be the first ever mutation experiments on declarative models. These experiments were conclusive in showing that the proposed mechanisms do not severely degrade the system performance in the chosen scenarios, yet usefully improve the system behavior in the presence of model faults.

There are many directions for future research. First, implementation of a plan analyzer should allow much better goal success levels to be achieved in the presence of faults since it should increase error detection coverage and provide lower latency. Implementation of the concurrent planning policy and comparison with the sequential planning policy are also of interest. Moreover, we would like to evaluate diversification on planning heuristics rather than just models and investigate also the additional detection capabilities of recent additions to the LAAS architecture [15]. Finally, many more experiments are needed to improve the statistical relevance of the results. The use of a large computer grid would drastically improve the number of experiments that could be executed in reasonable time and eliminate the need for manual inspection to remove trivial mutants.

References

[1] R. Alami, R. Chatila, S. Fleury, M. Ghallab, and F. Ingrand. An Architecture for Autonomy. *The International Journal of Robotics Research*, 17(4):315–337, April 1998.

[2] M. Becker and D. R. Smith. Model Validation in Planware. In *ICAPS 2005 Workshop on Verification and Validation of Model-Based Planning and Scheduling Systems*, Monterey, California, June 6-7, 2005.

[3] D. E. Bernard, E. B. Gamble, N. F. Rouquette, B. Smith, Y. W. Tung, N. Muscettola, G. A. Dorias, B. Kanefsky, J. Kurien, W. Millar, P. Nayal, K. Rajan, and W. Taylor. Remote Agent Experiment DS1 Technology Validation Report. Ames Research Center and JPL, 2000.

[4] I. R. Chen. On the Reliability of AI Planning Software in Real-Time Applications. *IEEE Transactions on Reliability*, 46(1):81–87, March 1997.

[5] I. R. Chen, F. B. Bastani, and T. W. Tsao. On the Reliability of AI Planning Software in Real-Time Applications. *IEEE Transactions on Knowledge and Data Engineering*, 7(1):14–25, February 1995.

[6] Y. Crouzet, H. Waeselynck, B. Lussier, and D. Powell. The SESAME Experience: from Assembly Languages to Declarative Models. In *Proceedings of the 2nd Workshop on Mutation Analysis (Mutation'2006)*, Raleigh, NC, November 7, 2006.

[7] M. Daran and P. Thvenod-Fosse. Software Error Analysis: a Real Case Study Involving Real Faults and Mutations. In *Proceedings of the 1996 ACM SIGSOFT international symposium on Software testing and analysis*, San Diego, California, January 8-10, 1996.

[8] A. E. Howe. Improving the Reliability of Artificial Intelligence Planning Systems by Analyzing their Failure Recovery. *IEEE Transactions on Knowledge and Data Engineering*, 7(1):14–25, February 1995.

[9] R. Howey, D. Long, and M. Fox. VAL: Automatic Plan Validation, Continuous Effects and Mixed Initiative Planning using PDDL. In *16th IEEE International Conference on Tools with Artificial Intelligence*, Boca Raton, Florida, November 15-17, 2004.

[10] S. Joyeux, A. Lampe, R. Alami, and S. Lacroix. Simulation in the LAAS Architecture. In *Proceedings of Principles and Practice of Software Development in Robotics (SDIR2005), ICRA workshop*, Barcelona, Spain, April 18, 2005.

[11] K. H. Kim and H. O. Welch. Distributed Execution of Recovery Blocks: An Approach for Uniform Treatment of Hardware and Software Faults in Real-Time Applications. *IEEE Transactions on Computers*, C-38:626–636, 1989.

[12] S. Lemai and F. Ingrand. Interleaving Temporal Planning and Execution in Robotics Domains. In *Proceedings of AAAI-04*, pages 617–622, San Jose, California, July 25-29, 2004.

[13] B. Lussier. *Fault Tolerance in Autonomous Systems*. PhD thesis, Institut National Polytechnique de Toulouse, 2007 (in French).

[14] B. Lussier, A. Lampe, R. Chatila, F. Ingrand, M. O. Killijian, and D. Powell. Fault Tolerance in Autonomous Systems: How and How Much? In *Proceedings of the 4th IARP/IEEE-RAS/EURON Joint Workshop on Technical Challenge for Dependable Robots in Human Environments*, Nagoya, Japan, June 16-18, 2005.

[15] F. Py and F. Ingrand. Real-Time Execution Control for Autonomous Systems. In *Proceedings of the 2nd European Congress ERTS, Embedded Real Time Software*, Toulouse, France, January 21-23, 2004.

[16] B. Randell. System Structure for Software Fault Tolerance. *IEEE Transactions on Software Engineering*, SE-1:220–232, 1975.