



SAPIENZA
UNIVERSITÀ DI ROMA

SmartPM: Automatic Adaptation of Dynamic Processes at Run-Time

Dipartimento di Ingegneria Informatica, Automatica e Gestionale A. Ruberti,
Università di Roma "La Sapienza"

Dottorato di Ricerca in Ingegneria Informatica – XXV Ciclo

Candidate

Andrea Marrella

ID number 796069

Thesis Committee

Prof. Massimo Mecella (Tutor)

Prof. Daniele Nardi

Prof. Luca Iocchi

Prof. Umberto Nanni

Prof. Stavros Vassos

Prof. Marco Aiello (Ext. Reviewer)

Prof. S. Rinderle Ma (Ext. Reviewer)

A thesis submitted in fulfillment of the requirements
for the degree of Doctor of Philosophy in Engineering in Computer Science

October 2013

Thesis defended on 7 October 2013
in front of a Board of Examiners composed by:

Prof. Andrea Clementi (chairman)
Prof. Paolo Ferragina
Prof. Roberto Basili

SmartPM: Automatic Adaptation of Dynamic Processes at Run-Time
Ph.D. thesis. Sapienza – University of Rome

© 2013 Andrea Marrella. All rights reserved

Website: <http://www.dis.uniroma1.it/~marrella>

Author's email: marrella@dis.uniroma1.it

Dedicated to my grandmother, with love.

Ringraziamenti

Questo lavoro di Tesi è il risultato di un percorso di ricerca durato 3 anni e mezzo. Provo emozioni contrastanti se ripenso all'esperienza di Dottorato che ho appena concluso. Ho vissuto momenti difficili, talvolta frustranti, ma l'impegno e la passione per la ricerca sono stati più forti delle difficoltà. In definitiva, non cambierei nulla della mia esperienza di Dottorato. Penso che mi abbia reso una persona migliore, dal punto di vista accademico ed umano.

Non sarei mai riuscito a completare questo percorso senza l'apporto di persone speciali. A questo proposito, vorrei principalmente ringraziare Chiara, che non ha mai mancato di incoraggiarmi e mi è stata vicina soprattutto nei momenti più difficili, e la mia famiglia, che mi ha sempre mostrato un sostegno impagabile in ogni situazione. Ringrazio anche i miei colleghi di Dottorato, in particolare Alessandro, Claudio, Francesco, Donatella, Mario, Riccardo, Paolo. Mi ritengo fortunato ad aver condiviso una parte della mia vita con loro.

Un ringraziamento speciale va a Massimo, grazie al quale ho acquisito un'esperienza lavorativa ed umana più unica che rara, e a quei ricercatori con cui ho lavorato e condiviso idee di ricerca. Una menzione particolare va al Prof. Yves Lesperance, grazie al quale ho potuto intraprendere una fantastica esperienza di ricerca e di vita in Canada, che mi ha arricchito culturalmente e professionalmente.

Grazie di cuore a tutti.

Contents

Extended Abstract	ix
1 Introduction	1
1.1 Flexibility Issues in Process Management Systems	2
1.2 The Spectrum of Process Management and Modeling Paradigms	7
1.2.1 Structured Processes	7
1.2.2 Loosely Structured Processes	8
1.2.3 Unstructured Processes	8
1.2.4 Dynamic Processes	9
1.3 Case Study	10
2 State of the Art	15
2.1 Process Adaptation	15
2.1.1 Exception Handling Techniques	16
2.1.2 Analysis of Existing PMSs	18
2.1.3 Discussion	20
2.2 AI-based Process Adaptation	21
2.2.1 Discussion	24
3 The SmartPM Approach	27
3.1 Overview of the Approach	27
3.1.1 Representing Tasks in SmartPM	29
3.1.2 Resource Model and Task Life-Cycle in SmartPM	30
3.2 Preliminaries	31
3.2.1 Situation Calculus	32
3.2.2 Indigolog	35
3.2.3 Classical Planning	39
3.3 Formalizing processes in IndiGolog	41
3.3.1 Realizing the Framework	53
3.4 Monitoring for Failures	55
3.5 The SmartPM Adaptation Mechanisms	60
3.5.1 The Built-in Adaptation Mechanism	60
3.5.2 The Plan-based Adaptation Approach	63
3.5.3 The Continuous Planning Approach	65
3.6 Conclusion	70

4	The SmartPM System	73
4.1	System Architecture	73
4.2	The IndiGolog Platform	76
4.2.1	The Top-level Main Cycle and Language Semantics	77
4.2.2	The Temporal Projector	79
4.2.3	The Communication Manager	79
4.2.4	The Domain Application	80
4.3	The SmartPM Definition Tool	81
4.3.1	The SmartML Modeling Language	81
4.3.2	Defining Processes in SmartPM through BPMN	87
4.3.3	The XML-to-IndiGolog Parser	89
4.4	Building the Planning Domain and the Planning Problem	92
4.5	SmartPM in Action	94
5	Validation	99
5.1	Performances of SmartPM in Computing Recovery Procedures	99
5.2	Effectiveness of SmartPM in Adapting Processes	100
6	Automatic Generation of Process Templates	109
6.1	Case Study	110
6.2	Partial-Order Planning	112
6.3	Process Templates	113
6.4	On Synthesizing a Library of Process Templates	118
6.4.1	The General Framework	119
6.4.2	Properties	122
6.5	Translation Algorithms	123
6.5.1	Representing Domain Theories and Process Cases in PDDL	123
6.5.2	Translating a Partially Ordered Plan P into a Process Template PT	125
6.6	Experiments	127
6.7	Related Work	130
6.8	Conclusion	131
7	Recovering Dynamic Processes in YAWL	133
7.1	Running Example	134
7.2	The YAWL Architecture	135
7.3	Making YAWL and SmartPM interoperate	136
7.4	The Planlets Approach	138
7.4.1	Incorporating Planlets into YAWL	140
7.4.2	Annotating YAWL Specifications in Planlets	142
7.5	Conclusion	150
8	Conclusion	153
A	The Full Code of the Example	155

Extended Abstract

Business Process Management [142] (a.k.a. BPM) is a “hot topic” because it is highly relevant from a practical point of view while at the same it offers many challenges for computer scientists and researchers. It is based on the observation that each product that a company provides to the market is the outcome of a number of activities performed. *Business processes* are the key instruments to organizing these activities and to improving the understanding of their interrelationships. BPM addresses the topic of process support in a broader perspective by incorporating different types of analysis (e.g., simulation, verification, and process mining) and linking processes to business and social aspects. Moreover, the current interest in BPM is fueled by technological developments (e.g., service oriented architectures) triggering standardization efforts. Several research issues have been addressed about the definition of models for describing processes, some of them more targeted towards domain-specific business designers (e.g., UML Activity Diagrams [35], BPMN – Business Process Modeling Notation [7]), and others more targeted to formal definitions of processes, in order to enable verification over process schemas (e.g., workflow nets [128] – a variant of Petri Nets [88, 98] targeted to describing processes, YAWL [124], etc.).

A *Process Management System* [130] (a.k.a. PMS) is a generic software system that is driven by explicit process representations (also called *process models*) to coordinate the enactment of business processes, aiming at increasing the efficiency and effectiveness in their execution. A process model, which is always built at design-time, is in charge of organizing the execution order of the activities of the business process. The basic constituents of a process model are *tasks* that describe an activity to be performed by an automated service (e.g., within a service-oriented architecture) or a human (e.g., an employee). The procedural rules to control the tasks are usually described by routing constructs like sequences, loops, parallel and alternative branches that form the *control flow* of the process [83].

The core of a PMS is the engine that takes in input a process model and manages the process routing by deciding which tasks are enabled for execution, taking into account the control flow, the value of process variables and tasks constraints. The representation of a single execution of a process model within the engine of the PMS is called a *process instance* [55]. Once a task is ready for being assigned, the engine is also in charge of assigning it to proper participants; this step is performed by considering the participant “skills” required by the single task: a task will be assigned to the participant that provides all of the skills required. Participants are provided with a client application, part of the PMS, named *Task Handler*. It is aimed at receiving notifications of task assignments. Participants can, then, use

this application to pick the next task to work on. Current technologies exist on the market which concretely allow the enactment of processes (e.g., the YAWL Engine [124] and the jBPM orchestration engine [22] based on the WS-BPEL [92] specification, etc.).

Traditionally, PMSs have focused on the support of predictable and repetitive business processes, which can be fully pre-specified in terms of formal process models. All possible paths through those processes are well-understood, and the process participants usually do not need to make a decision about what to do next since the path is completely determined by their data entry or other attributes of the process. This kind of highly-structured work includes mainly production and administrative processes (such as financial services, manufacturing, etc.) [73]. However, current maturity of process management methodologies has led to the application of process-oriented approaches in new challenging knowledge-intensive scenarios [33], such as healthcare [102, 69] or home automation (e.g., domotics [53]). In such working environments, changes in the operational context and in other heterogeneous contextual information may occur unpredictably and at any time, requiring the ability to react to those changes and properly *adapt* and *modify* process behavior. This has led to the need to provide support for flexible and adaptive process management [141, 105], by reconsidering the trade-off between flexibility and support provided by existing PMSs [135]. According to [115], *Process Adaptation* can be seen as the ability of a process to react to exceptional circumstances (that may or may not be foreseen) and to adapt/modify its structure accordingly.

The current-day leading commercial PMS products [20, 125, 60, 85] and research prototypes [17, 18, 124] provide some techniques to react to exceptions and adapt process instances to mitigate their effects. Specifically, they provide the support for the handling of *expected exceptions*. The process models are designed in order to cope with potential exceptions, i.e., for each kind of exception that is envisioned to occur, a specific contingency process (a.k.a. *exception handler* or compensation flow) is defined. These approaches perform well with predictable and repetitive business processes, where all possible exceptions and deviations that can be encountered are predictable and defined in advance, along with the specific handling logic. However, in knowledge-intensive scenarios, the process usually *dynamically evolves*, i.e., it strongly depends on user decisions made during process execution. For those processes, it is not possible to anticipate all real-world exceptions and to capture their handling in a process model at design-time.

In this direction, since the last Nineties, a new class of *Adaptive Process Management Systems* is emerged, by facilitating structural changes of processes at run-time [103, 48, 87, 104, 109, 139]. When something goes wrong during the process execution, structural changes apply directly to process elements, and the adaptation is carried out by deleting, adding, or modifying one or several process elements. For example, an adaptive PMS like ADEPT2 [104] is able to support the handling of unanticipated exceptions, by enabling different kinds of ad-hoc deviations from the pre-modeled process instance at run-time, according to the structural process change patterns defined in [138]. New process models can be created and tailored for a particular demand or business case, and process instances can be adapted after they have been started if some unforeseen events occur. Currently, adaptive PMSs have reached such a level of maturity that they are about to be transferred into

practice [83].

The majority of the above approaches face the challenge to provide flexibility and adaptation and to offer process support at the same time. Traditional PMSs deal with expected exceptions at design-time by automatically providing an exception handler at run-time, while adaptive PMSs offers structural process change at run-time for unanticipated exceptions, but they do not automate the adaptation; a manual intervention of a domain expert is always required for adapting a faulty process instance at run-time. However, in the last years, the widespread availability of mobile computing platforms has led to the application of process-oriented approaches in *pervasive and highly dynamic scenarios* [16, 31, 14, 15]. An interesting example comes from the emergency management domain. During the management of complex emergency scenarios, teams of first responders act in disaster locations with the main purpose of achieving specific goals, including assisting potential victims and assessing and stabilizing the situation. The set of activities and procedures that collectively define an emergency response plan are characterized for being as complex as typical business processes and for involving teams of many members. Emergency response operators can benefit from the use of mobile devices and wireless communication technologies, as well as from the adoption of a process-oriented approach for team coordination [114]. A response plan encoded as a business process and executed by a PMS deployed on mobile devices can help coordinate the activities of emergency operators equipped with PDAs and smartphones and supported by mobile networks. In this dynamic context the environment may change continuously and processes can be easily invalidated because of exogenous events and of tasks not executed as expected. This means that (i) it is not possible to predict all possible exceptions at design-time and that (ii) adaptation ought to be as *automatic* as possible and to require *minimum manual human intervention* at run-time; in fact, in emergency management, saving minutes could result in saving injured people, preventing buildings from collapses, and so on.

Research Contributions

The main focus of the author’s research activity is to devise an approach for *run-time automatic adaptation of dynamic processes*. Dynamic processes are a particular kind of processes for which there is not a clear, anticipated correlation between a change in the context and corresponding process changes. Usually, the structure of a dynamic process can be completely captured with a procedural process model that explicitly defines the tasks and their execution constraints. Examples of dynamic processes are processes for emergency management (an extensive case study is presented in Section 1.3) and military forces deployment plans. A dynamic process is thought to be enacted in pervasive and highly dynamic scenarios, where exceptions and exogenous events “are not the exception but the rule”. Dynamic context changes or undesirable outcome of some activities may often cause abnormal termination of the process activities and prevent the achievement of the business goals.

Traditional approaches that try to anticipate how the work will happen by solving each problem at design time [20, 125, 60, 85, 17, 18, 124], as well as approaches that allow to manually change the process structure at run time [103, 48, 87, 104, 109, 139],

are often ineffective or not applicable in rapidly evolving contexts. The design-time specification of all possible compensation actions requires an extensive manual effort for the process designer, that has to anticipate all potential problems and ways to overcome them in advance, in an attempt to deal with the unpredictable nature of dynamic processes. Moreover, the designer often lacks the needed knowledge to model all the possible contingencies, or this knowledge can become obsolete as process instances are executed and evolve, by making useless his/her initial effort. Although the exploitation of current adaptive PMSs to support the enactment of processes in pervasive and mobile scenarios represents a promising and helpful approach, dynamic processes demand a more agile approach recognizing the fact that in dynamic environments process models quickly become outdated and hence require closer interweaving of modeling and execution [105].

With respect to these needs, the author's research main target is to devise *intelligent failure handling mechanisms* that allow to monitor running process instances and to react to tasks failures or to the occurrence of exogenous events that may put at risk process executions, by providing *automatic adaptation* to identify a suitable compensation and recovery strategy. The idea is to provide an adaptation framework that is able to deal automatically with *unanticipated exceptions at run-time*, without explicitly defining any handler/policy to recover from exceptions at design-time, and without the intervention of domain experts.

If compared with classical business processes, the execution and adaptation of dynamic processes demand two special requirements:

- A dynamic process must be *adaptable to the context*. This means there is the need to explicitly define the contextual data describing the scenario in which the process will be enacted;
- A PMS that executes a dynamic process must be able to *automatically detect* exceptional situations, to *derive* and to correctly *apply* the recovery procedures necessary to handle them.

To this end, we use a specialized version of the concept of adaptation from the field of agent-oriented programming [27]. Specifically, we consider adaptation as reducing the gap between the *expected reality*, the (idealized) model of reality that is used by a PMS to reason on the dynamic process under execution, and the *physical reality*, the real world with the actual values of conditions and outcomes. A misalignment of the two realities often stems from errors in the tasks outcomes (e.g., incorrect data values) or is the result of exogenous events coming from the environment.

Our approach to process adaptation is mainly based on well-established techniques and frameworks from Artificial Intelligence (a.k.a. AI), such as *Situation Calculus*[107], *IndiGolog* [25] and *automatic planning* [90]. Situation Calculus is a logical language specifically designed for representing dynamically changing worlds in which all changes are the result of named actions. We used Situation Calculus for providing a declarative specification of the domain (i.e., available tasks, contextual properties, tasks preconditions and effects, what is known about the initial state) in which the dynamic process has to be executed. On top of Situation Calculus, we used the *IndiGolog* language for formalizing the structure and the control flow of our dynamic process. *IndiGolog* is a logic-based programming language used for robot and agent programming.

Then, we propose a PMS realization, namely SmartPM [30, 29], which is based on the IndiGolog interpreter¹ developed at University of Toronto and RMIT University, Melbourne. The IndiGolog interpreter reasons about the preconditions and effects of the process tasks to find a legal terminating execution of the process. IndiGolog programs are executed online together with sensing the environment and monitoring for events. When an exception or an exogenous event is sensed, it results in a discrepancy between the physical and expected reality, and the IndiGolog interpreter is in charge to determine if such event is able to invalidate the execution of the dynamic process under execution. If so, SmartPM allows the synthesis of a recovery procedure at run-time by invoking an external *state-space planner* [45]. In AI, planning systems are problem-solving algorithms that operate on explicit representations of an initial state (the faulty process state that reflects the physical reality), a goal condition (the process state reflecting the expected reality) and actions (the set of tasks executable in the contextual scenario under observation). A state-space planner explores only strictly linear sequences of actions directly connected from the initial state to the goal, i.e., in our case, it searches for a plan that may turn the physical reality into the expected reality. If the recovery plan exists, it will be executed by SmartPM for adapting the faulty process instance. Since the adaptation mechanism deployed on SmartPM is *blocking* (i.e., the execution of the main process is stopped during the synthesis/enactment of the recovery procedure), we also propose a *non-blocking* repairing mechanism based on *continuous planning* techniques [77], that avoids to stop directly any task in the main process during the computation of the recovery process.

In the second part of the thesis, we analyze a problem that often involves the design-time specification of a dynamic process. Since resources of a dynamic process are usually shared by the process participants, it is difficult to foresee all the potential tasks interactions in advance and there is the risk that concurrent tasks could not be independent from one from another (e.g., they could operate on the same data at the same time), resulting in incorrect outcomes. We address this issue and proposing an approach [76] that exploits *partial-order planning algorithms* [90, 140] for building automatically a library of process template definitions. Partial-order planning differs from classical state-space planning algorithms, that explore only strictly linear sequences of actions directly connected to the start or goal, by devising totally ordered plans. On the contrary, partial-order planning is based on the least commitment principle [140], whose main advantage is that decisions about action ordering are postponed until a decision is forced, thus guaranteeing flexibility in the execution of the plan and by possibly permitting actions to run concurrently. The strength of the approach is that resulting templates are reusable in a variety of partially-known contextual environments, and all concurrent tasks composing the templates are effectively independent one from another.

Finally, although SmartPM is born as a PMS for supporting first responders in emergency management scenarios, the planning-based adaptation approach employed is enough general for being used on top of existing PMSs. Specifically, in the third part of this Thesis, we concretize our approach to automatic process adaptation on

¹<http://sourceforge.net/projects/indigolog/>

top of the well-known YAWL modeling language and execution environment [124]. To this end, we show a concrete design and implementation proposal of how the YAWL architecture can be extended to integrate planning capabilities and to support the handling of unanticipated exception at run-time [80, 79, 81].

In summary, the main contributions and results of the author's research activity can be summarized as follow:

- R1** First of all, in [28] we collected and analyzed the general requirements for the application of a process-oriented approach in emergency management scenarios. Starting from those requirements, in [31] we focus on the challenges related to the design and implementation of a PMS able to support first responders on the field, analyzing the core components of the overall architectural model.
- R2** In [16] we present the high level architecture for emergency management systems devised in the European research project WORKPAD². Such an architecture is specifically tailored in supporting collaborative work of human operators during disaster scenarios, where different teams, belonging to different organizations, need to collaborate in order to reach a common goal. The overall approach, user-centered methodology and achieved results in the design, implementation and validation of the WORKPAD architecture are detailed in [57, 58, 10, 15]. Moreover, a qualitative and quantitative evaluation of main outcomes, as a result of on-the-field validation and showcase activities of the project are shown in [14].
- R3** Starting from the experience gained in the area and lessons learned, in [78] we provide a general set of guidelines, suggestions and possible research directions on how to effectively design mobile information systems for supporting on-the-field collaboration of emergency operators.
- R4** In [21] we present the design and prototype of a PMS for improving operational support to clinicians during their daily activities in hospital wards, on the basis of clinical guidelines.
- R5** In [33], starting from three different dynamic real world scenarios, we present a critical and comparative analysis of the existing approaches used for supporting, modeling and adapting dynamic and knowledge-intensive processes.
- R6** In [30, 29, 77] we present and formalize SmartPM, our AI-based PMS that deals automatically with unanticipated exceptions at run-time. Specifically, in [30, 29] we exploit a built-in adaptation mechanism offered by the IndiGolog interpreter (basically, a simple planner based on a breadth-first search algorithm) for the synthesis of the recovery procedure. In order to improve drastically the time needed for finding the recovery plan, in [77] we delegate every aspect of adaptation to an external state-of-the-art planner. This allows to separate the planning phase with the modeling and execution phase, and to introduce a *non-blocking* repairing mechanisms, based on *continuous planning* techniques.
- R7** In [80] we contextualize and demonstrate our approach in a service-oriented environment, as an application of the architectural solutions for the integration of different modeling approaches to achieve flexibility [129]. We show

²<http://www.dis.uniroma1.it/workpad/>

how the YAWL environment (and its imperative modeling approach) can be complemented with the SmartPM execution environment [29].

- R8** In [79] we propose a general approach and a conceptual architecture to automatic process adaptation, based on the concept of declarative modeling of processes and the use of continuous planning techniques; we show the feasibility of the proposed approach by discussing its deployment on top of the YAWL system [124].
- R9** In [81], we introduce and define PLANLETS, as self-contained YAWL specifications where process tasks are annotated at design-time with pre-conditions, desired effects and post-conditions. The role of declarative task annotations is twofold: (i) pre- and post-conditions enable run-time process execution monitoring and exception detection: they are checked respectively before and after task executions, and the violation of a pre- or post-condition results in an exception to be handled; (ii) along with the input/output parameters consumed/produced by the task, pre-conditions and effects provide a complete specification of the task: this allows the task to be represented as an action in a planning domain description and used for solving a planning problem built to handle an exception. In the presence of an exception, this approach allows delegating to an external planner the automatic run-time synthesis of a suitable recovery procedure by contextually selecting the compensation tasks from a specific repository linked to the PLANLET under execution.
- R10** Since the design time specification of dynamic processes can be time-consuming and error-prone, due to the high number of tasks involved and their context-dependent nature, such processes frequently suffer from potential interference among their constituents, since resources are usually shared by the process participants and it is difficult to foresee all the potential tasks interactions in advance. Concurrent tasks may not be independent from one from another (e.g., they could operate on the same data at the same time), resulting in incorrect outcomes. To address these issues, in [76] we propose an approach that exploits partial-order planning algorithms for automatically synthesizing a library of process template definitions for different contextual cases. The resulting templates guarantee sound concurrency in the execution of their activities and are reusable in a variety of partially-known contextual environments.

During the realization of this thesis, the following publications have been produced:

- [76] **A. Marrella, Y. Lespérance**
Synthesizing a Library of Process Templates through Partial-Order Planning Algorithms. 14th International Working Conference on Business Process Modeling, Development and Support (BPMDS 2013), Valencia, Spain, 17-18 June 2013.
- [81] **A. Marrella, A. Russo, M. Mecella**
Planlets: Automatically Recovering Dynamic Processes in YAWL. 20th International Conference on Cooperative Information Systems (CoopIS 2012) - OTM Conferences (1), Rome, Italy, 10-14 September 2012.
- [21] **F. Cossu, A. Marrella, M. Mecella, A. Russo, G. Bertazzoni, M. Suppa, F. Grasso**
Improving Operational Support in Hospital Wards through Vocal Interfaces and Process-Awareness. 25th IEEE International Symposium on Computer-Based Medical Systems (CBMS 2012), Rome, Italy, 20-22 June 2012.

- [33] **C. Di Ciccio, A. Marrella, A. Russo**
Knowledge-intensive Processes: An Overview of Contemporary Approaches. 1st International Workshop on Knowledge-intensive Business Processes (KiBP 2012), Rome, Italy, 15 June 2012.
- [79] **A. Marrella, M. Mecella, A. Russo**
Featuring Automatic Adaptivity through Workflow Enactment and Planning. 7th International Conference on Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom 2011), Orlando, Florida, USA, 15-18 October 2011.
- [80] **A. Marrella, M. Mecella, A. Russo, A.H.M. ter Hofstede, S. Sardina**
Making YAWL and SmartPM Interoperate: Managing Highly Dynamic Processes by Exploiting Automatic Adaptation Features. 9th International Conference on Business Process Management (BPM 2011), Demonstration Track, Clermont-Ferrand, France, 28 August - 02 September 2011.
- [29] **M. de Leoni, A. Marrella, M. Mecella, S. Sardina**
SmartPM - Featuring Automatic Adaptation to Unplanned Exceptions. Technical Report of Dipartimento di Informatica e Sistemistica ANTONIO RUBERTI, SAPIENZA - Università di Roma. June 2011.
- [77] **A. Marrella, M. Mecella**
Continuous Planning for solving Business Process Adaptivity. 12th International Working Conference on Business Process Modeling, Development and Support (BPMDS 2011), London, UK, 20-21 June 2011.
- [15] **T. Catarci, M. de Leoni, A. Marrella, M. Mecella, A. Russo, M. Bortenschlager, R. Steinmann**
WORKPAD : Process Management and Geo-Collaboration Help Disaster Response. International Journal of Information Systems for Crisis Response and Management (IJISCRAM), Volume 3, Issue 1, pp. 32–49, 2011.
- [78] **A. Marrella, M. Mecella, A. Russo**
Collaboration On-the-field: Suggestions and Beyond. 8th International Conference on Information Systems for Crisis Response and Management (ISCRAM 2011), Lisbon, Portugal, 8-11 May 2011.
- [31] **M. de Leoni, A. Marrella, A. Russo**
Process-aware Information Systems for Emergency Management. International Workshop on Emergency Management through Service Oriented Architectures (EMSOA) co-located with the ServiceWave 2010 Conference, Ghent, Belgium, 13 December 2010.
- [14] **T. Catarci, M. de Leoni, A. Marrella, M. Mecella, M. Bortenschlager, R. Steinmann**
The WORKPAD Project Experience: Improving the Disaster Response through Process Management and Geo Collaboration. 7th International Conference on Information Systems for Crisis Response and Management (ISCRAM 2010), Seattle, USA, 2-5 May 2010.
- [58] **S. R. Humayoun, T. Catarci, M. de Leoni, A. Marrella, M. Mecella, M. Bortenschlager, R. Steinmann**
The WORKPAD User Interface and Methodology: Developing Smart and Effective Mobile Applications for Emergency Operators. 13th International Conference on Human-Computer Interaction (HCI International 2009), Session “Designing for Mobile Computing”, San Diego, USA, 19-24 July 2009.
- [57] **S. R. Humayoun, T. Catarci, M. de Leoni, A. Marrella, M. Mecella, M. Bortenschlager, R. Steinmann**
Designing Mobile Systems in Highly Dynamic Scenarios. The WORKPAD Methodology. Springer’s International Journal on Knowledge, Technology and Policy, Volume 22, Number 1 - March 2009.

- [30] **M. de Leoni, A. Marrella, M. Mecella, S. Valentini, S. Sardina**
Coordinating Mobile Actors in Pervasive and Mobile Scenarios: An AI-based Approach. 2nd IEEE International Workshop on Interdisciplinary Aspects of Coordination Applied to Pervasive Environments: Models and Applications (COMA 2008) at WETICE 08, Rome, Italy, 23-25 June 2008.
- [10] **A. Capata, A. Marrella, R. Russo, M. Bortenschlager, H. Rieser**
A Geo-based Application for the Management of Mobile Actors during Crisis Situations. 5th International Conference on Information Systems for Crisis Response and Management (ISCRAM 2008), Washington DC, USA, 4-7 May 2008.
- [16] **T. Catarci, M. de Leoni, A. Marrella, M. Mecella, B. Salvatore, G. Vetere, S. Dustdar, L. Juszczak, A. Manzoor, Hong-Linh Truong**
Pervasive and Peer-to-Peer Software Environments for Supporting Disaster Responses. IEEE Internet Computing Journal - Special Issue on Crisis Management - Volume 12, Number 1 - January-February 2008.
- [28] **M. de Leoni, A. Marrella, M. Mecella, F. De Rosa, A. Poggi, A. Krek, F. Manti**
Emergency Management: from User Requirements to a Flexible P2P Architecture. 4th International Conference on Information Systems for Crisis Response and Management (ISCRAM 2007), Delft, the Netherlands, 13-16 May 2007.

The work [76], presented thoroughly in Chapter 6, is the result of a research internship of the author at the Department of Computer Science and Engineering at York University in Toronto (Ontario, Canada), under the supervision of Prof. Yves Lespérance.

The implementation of the IndiGolog based PMS has been developed in cooperation with Dr. Sebastian Sardina, research fellow at the Intelligent Agent Group of the RMIT University in Melbourne (Australia) and Dr. Massimiliano de Leoni, research assistant at the Eindhoven University of Technology (the Netherlands).

The works [80] and [81] are the result of a research collaboration with Prof. Arthur H. M. ter Hofstede, the co-leader of the BPM Group of the Faculty of Information Technology of Queensland University of Technology, Brisbane (Australia).

The author has also co-chaired a workshop on Knowledge-intensive Business Processes (KiBP 2012) held in Rome on June 15th, 2012, co-located with the 13th International Conference on Principles of Knowledge Representation and Reasoning (KR 2012)³. The main focus of the workshop was to discuss about how the use of techniques that came from different fields, such as Artificial Intelligence (AI), Knowledge Representation (KR), Business Process Management (BPM), Service Oriented Computing (SOC), etc., can be used jointly for improving the modeling and the enactment phase of a knowledge-intensive process. The purpose was to devise interesting approaches that can still achieve the goals of understanding, visibility and control of these emergent processes. The KiBP 2012 proceedings are available online at <http://ceur-ws.org/Vol-861/>.

Thesis Outline

- Chapter 1 introduces background concepts and definitions related to process flexibility and process adaptation, and provides a systematic view of the

³Web site: <http://www.dis.uniroma1.it/~kibp2012/>

different approaches and methodologies that have emerged to support classes of processes with different requirements. This serves as the basis for positioning the performed work. Moreover, a case study based on a real emergency management scenario is presented.

- Chapter 2 analyzes the state of the art concerning process adaptation and process recovery. Specifically, we first analyze existing process adaptation and exception handling techniques. Then, we discuss on the degree of adaptation/flexibility provided by several commercial PMSs and academic prototypes. Finally, we directly compare **SmartPM** with other research works that exploit AI techniques for improving the degree of process adaptation.
- Chapter 3 focuses on our general approach to automatic process adaptation, that involves formalizing processes in Situation Calculus and **IndiGolog**. We clearly define the execution semantic provided by **SmartPM** for the enactment of dynamic processes, together with the logic used for monitoring running process instances and adapting them when needed through the use of classical planning techniques.
- Chapter 4 discusses the overall architecture of **SmartPM**, our AI-based PMS, including technical details about the system implementation, the **IndiGolog** interpreter and the Task Handler used for assigning tasks to process participants. Then, we describe the **SmartPM Modeling Language** (a.k.a. **SmartML**), which combines a modeling formalism for representing the information of the contextual scenario linked to a specific dynamic process, and a graphical tool (specifically, Eclipse BPMN⁴) for designing the control flow of the process. We also show how a dynamic process formalized through **SmartML** is automatically translatable in Situation Calculus and **IndiGolog** readable formats and is therefore ready for being executed by **SmartPM**. Finally, we analyze the algorithms for converting a **SmartML** domain theory in a PDDL planning domain and for generating a PDDL planning problem when process adaptation is required.
- Chapter 5 reports on performance evaluation and system validation activities. Specifically, we first report on experimental evaluation results, in terms of time needed for automatically adapting the dynamic process taken from our case study when exceptions of growing complexity arise. Then we measure the effectiveness of **SmartPM** in finding recovery procedures by simulating the execution of thousands of processes instances having different control-flows in different contextual environments.
- Chapter 6 proposes an approach that exploits partial-order planning algorithms for synthesizing automatically a library of process template definitions starting from a declarative specification of process tasks. A template can be seen as the “closest thing” to a completely defined process model. It guarantees sound concurrency in the execution of its activities and is reusable in a variety of partially-known contextual environments.
- Chapter 7 provides an in-depth discussion, a concrete design and implementation proposal of how the YAWL architecture can be extended to integrate planning capabilities. For this aim, we propose the approach of **PLANLETS**,

⁴<http://www.eclipse.org/modeling/mdt/?project=bpmn2>

self-contained YAWL specifications with recovery features, based on modeling of pre- and post-conditions of tasks and the use of planning techniques.

- Chapter 8 concludes the thesis by discussing limitations and future developments of the planning-based adaptation approach we have proposed. Moreover, we show ongoing and future research activities we are currently investigating.

Chapter 1

Introduction

Business process management (BPM) solutions have been prevalent in both industry products and academic prototypes since the late 1990s [75]. A classical *business process* reflects a “preferred work practice”, i.e., a set of one or more connected activities which collectively realize a particular business goal [142]. Usually, a business process is linked to an organizational structure defining functional roles and organizational relationships. Examples of business processes include insurance claim processing, order handling and personnel recruitment [105].

In order to improve their business processes, enterprises are increasingly interested in aligning their information systems in a process-centered way offering the right business functions to the right users at the right time [36, 105]. This need has evolved primarily from the desire to understand, organize, and automate the processes upon which a business is based. For this purpose, during the last decade, a new generation of information systems, called *Process Management Systems* (a.k.a. PMSs, or more generally, Process-Aware Information Systems, a.k.a. PAISs) have become increasingly popular to effectively support the business processes of a company at an operational level. A PMS is a software system that manages and executes operational processes involving people, applications and information sources on the basis of process models [36].

PMSs hold the promise of facilitating the everyday operation of many enterprises and work environments, by supporting business processes in all the steps of their life-cycle [36]. As shown in Figure 1.1(a), the “life” of a business process is organized in 4 main stages. In the *design* phase, starting from a requirements analysis, *process models* are designed using a suitable modeling language. A process model is a formal representation of work procedures that controls the sequence of performed tasks and the allocation of resources to them [93]. In the *configuration* phase process models are implemented by configuring a PMS in order to support process enactment via an execution engine. In the *enactment* phase *process instances* are then initiated, executed and monitored by the run-time environment. The execution engine drives and monitors the work of the involved entities, and performed tasks generating execution traces are tracked and logged. After process execution, in the *diagnosis* phase process logs are evaluated and mined to identify problems and possible improvements, potentially resulting in process re-design and evolution.

Traditionally, PMSs have focused on the management of “administrative” pro-

cesses characterized by clear and well-defined structures. However, the use of business processes for supporting work in highly dynamic contexts (such as emergency management or health care) has become a reality, thanks also to the growing use of mobile devices in everyday life, which offer a simple way of picking up and executing tasks. Those processes are usually subject to an higher frequency of unexpected contingencies than classical scenarios; therefore, a certain degree of flexibility is needed to support dynamic process adaptation in case of exceptions.

In this chapter, that serves as the basis for positioning the performed author's research activities, we aim at discussing the flexibility requirements of actual business processes, ranging from pre-specified processes to dynamic and knowledge-intensive processes. For this purpose, in Section 1.1 we analyze concepts and definitions related to process flexibility and in Section 1.2 we classify business processes on the basis of their degree of structure by providing a systematic view of the different approaches and methodologies that have emerged to support classes of processes with different requirements. Specifically, we will focus on addressing fundamental adaptation needs for supporting dynamic processes in mobile settings for pervasive scenarios. To this end, in Section 1.3 we present a case study based on a real dynamic process thought to be enacted in emergency management scenarios.

1.1 Flexibility Issues in Process Management Systems

The notion of flexibility has emerged as a main research topic in BPM over the last years [121, 59, 116]. The need for flexibility stems from the observation that organizations often face continuous and unprecedented changes in their respective business environments. Such disturbances and perturbations of business routines need to be reflected within the business processes in the sense that processes need to be able to adapt to such change. Research on process flexibility has traditionally explored alternative ways of considering flexibility during the design of a business process. The focus typically has been on ways of how the demand for process flexibility can be satisfied by advanced process modeling techniques, i.e., issues intrinsic to the processes [110].

In [115], the authors define flexibility as “the ability of the process to execute on the basis of a loosely, or partially specified model, where the full specification of the model is made at runtime, and may be unique to each instance”. Moreover, they advocate an approach that aims at making the process of change part of the business process itself. Such an approach, called *Pockets of Flexibility*, relies on a pre-specified process model with placeholder activities. For each placeholder activity, a constraint-based process model (i.e., activities and constraints) can be specified. During process enactment, placeholder activities are refined, meaning that users define a process fragment that has to satisfy the constraints and which substitutes the placeholder activity.

An interesting look into the various ways in which flexibility can be achieved is made in [117]. Here an extensive taxonomy of process flexibility is proposed. In particular, it is presented a comprehensive description of four distinct approaches that can be taken to facilitate flexibility within a process. *Flexibility by Design* is the ability to incorporate alternative execution paths within a process model at design




Figure 1.1. Process life-cycle (a) and Spectrum (b).

time. The selection of the most appropriate execution path is made at run-time for each process instance. *Flexibility by Deviation* is presented as the ability for a process instance to deviate at run-time from the execution path prescribed by the original process without altering its process model. *Flexibility by Underspecification* is the ability to execute an incomplete process model at run-time, i.e., the model needs to be completed by providing a concrete realization for the undefined parts at run-time. Finally, *Flexibility by Change* is the ability to modify a process model at run-time such that one or all of the currently executing process instances are migrated to a new process model. All of these strategies try to improve the ability of business processes to respond to changes in their operating environment without necessitating a complete redesign of the underlying process model, however they differ in the timing and manner in which they are applied. Moreover they are intended to operate independently of each other.

The above works concentrated on the control-flow perspective of a business process, while other perspectives addressing data and resources used in a process are also subject to flexibility requirements. A complete analysis that incorporates also these perspectives has been performed in [105], where flexible process support is characterized with four major flexibility needs, namely support for (i) *variability*, (ii) *looseness*, (iii) *adaptation*, and (iv) *evolution*.

- *Process variability* requires processes to be handled differently - resulting in different *process variants* - on the basis of a given context. Starting from a fixed core process model, the course of actions may vary from variant to variant [49]. Usually, there exists a multitude of variants of a particular process model, whereby each of these variants is valid in a specific scenario; i.e., the configuration of a particular process variant depends on requirements of the process context. Variability can be usually introduced due to different groups of customers involved in the process enactment or differences in regulations found in different countries [50].

Example 1.1.1. A typical example of a real procedure that needs flexible support is the process for the organization of the study plans for master students in some European universities. The procedure for application, review and acceptance of study plans is managed by an on-line system and is generally well defined. After the enrollment, all the master students must submit a study plan from the on-line system. Let us suppose that the courses for the study plan can be chosen starting from 3 specialization options (Computer Networks, Software and Services for the Information Society, Distributed Systems and Architectures) with pre-selected combination of courses. If a student chooses to include in its study plan a pre-selected set of courses, the approval of the study plan is immediate (and the review phase is not required).

Variability can be introduced if a student decide to choose freely which courses to include in the study plan, by combining the single courses taken from each of the specialization. In this last case, the study plan must be reviewed by a university delegate, modified (if necessary), and, finally, approved or rejected. Another variant of the procedure is enacted when a foreign student wants to apply for a study plan. These kinds of students can not apply their study plan on-line, but they need to be received by the university delegate, and the study plan is built ad-hoc for the single cases.

- *Looseness* is a characteristic of *Knowledge-intensive Processes* [33], that are processes characterized by being *non repeatable* (the models of two process instances may differ one another), *non predictable* (the course of the actions depends on context-specific parameters, whose values are not known a priori and may change during process execution) and *emergent* (the course of the actions only emerges during process execution, when more information is available) [105]. For those processes, only the goal and the modeling of the loose process are known a priori, meaning that those processes can not be fully pre-specified at design-time.

Example 1.1.2. Typically, for being enacted, health-care processes requires a loose specification. For example, Patient Treatment Processes are rarely identical and the course of actions is unpredictable, since it depends on the specific patient's case (e.g., health status of the patient, allergies, examination results, etc.).

- *Process Adaptation* is the ability of a process to react to exceptional circumstances (that may or may not be foreseen) and to adapt/modify its structure accordingly [115]. Exceptions can be either *expected* or *unanticipated* [105]. An expected exception can be planned at design-time, i.e., a process designer can provide an *exception handler* which is invoked during run-time to cope with the expected exception. Therefore, if during run-time the PMS detects an expected exception, it immediately invokes a suitable exception handler for dealing with the exception itself. A list of *exception handling patterns* that can be applied when exceptional situations arise is shown in [71]. Those patterns cover typical strategies to be used when defining exception handlers for a particular process model.

Example 1.1.3. If we consider the procedure for managing study plans, a “classical” expected exception is captured when a student forgets to compile all the mandatory fields (e.g., the student ID code or the mobile phone number) needed for submitting the study plan on-line. In such a case, the on-line system notifies that some required information is missing, and asks to the student to compile correctly all the mandatory fields. The system will allow to submit the study plan only when all required information has been correctly provided.

Even though the handling of expected exceptions is fundamental for every PMS, in many cases the number of possible exceptions may be too large, and requires an extensive manual effort for the process designer, that has to anticipate all potential problems and ways to overcome them in advance. Moreover, for many knowledge-intensive and dynamic scenarios, expected exceptions cover only partially relevant situations [123], and it is not realistic to assume that all exceptional situations, as well as required exception handlers, can be anticipated at design-time and thus incorporated into the process model [105]. This means that PMSs should provide the support for the handling of *unanticipated exceptions* at run-time. Such exceptions can be detected during the execution of a process instance, when a mismatch between the computerized version of the process and the corresponding real-world business process occurs. To cope with those exceptions, a PMS is required to allow *structural adaptation* of its corresponding process model at run-time; structural changes apply directly to process elements, and the adaptation is carried out by deleting, adding, or modifying one or several process elements.

Example 1.1.4. In many universities, it may happen that a student passes a course that is not listed in the her/his study plan. Such an exception can not be managed at design-time, since the university delegate, which supervises the procedure for managing study plans, may decide:

- to include the passed course in the student’s study plan as an “excess course”;
- to insert the passed course in the student’s study plan by substituting it with another similar course already included in the study plan;
- to refuse to insert the passed course in the student’s study plan.

Each decision requires a certain knowledge about the student’s academic situation. For example, if a student of Computer Science Engineering passes an exam of Anatomy at the Faculty of Medicine, it is obvious that the university delegate will refuse to update the student’s study plan by inserting the respective course (that is out of the scope of the student’s course of studies). On the contrary, if a student passes a exam during a period spent in a foreign university, and the associate course is not included in her/his study plan but is very similar to another course included in her/his study plan, the university delegate may decide to substitute the planned course with the one passed by the student when s/he was abroad.

- *Process Evolution* is the ability of an implemented process to change when the corresponding business process evolves [12]. It is often driven by changes

in the business, the technological environment, and the legal context [132]. The evolution may be *incremental* as for process improvements [23] (i.e., only small changes are required to the implemented process), or *drastic* as for process innovation or process re-engineering [52] (i.e., if radical changes are required). The biggest problem here concerns the handling of active process instances, which were initiated in an old model, but need to comply with a new specification. Achieving compliance for these affected instances may involve loss of work and therefore has to be carefully planned [115].

Example 1.1.5. Let us consider again the procedure for managing study plans, and suppose that the University revises the admission procedure for master students, requiring all applicants to submit a statement of purpose together with their application for the study plan. To implement this change, there can be two options available; one is to flush all existing applications, and apply the change to new applications only. Thus all existing applications will continue to be processed according to the old process model. The second option to implement the change is to migrate to the new process. It may be decided that all applicants, existing and new, will be affected by the change. Thus all admission applications, which were initiated under the old rules, now have to migrate to the new process.

With respect to these needs, our research activities have broadly focused on the problem of *process adaptation* defined as the ability of a PMS to adapt the process and its structure (i.e., pre-specified model) to emerging events. While several approaches [20, 125, 60, 85, 17, 18, 124] have been proposed and implemented for dealing with *expected* exceptions via exception handlers typically pre-specified by process designers at design-time, we tackled the problem of *unanticipated* exceptions and their handling through structural process changes at run-time. Specifically, we focus our attention on *dynamic processes*, a particular kind of processes thought to be enacted in pervasive and highly dynamic scenarios, and for which there is not a clear, anticipated correlation between a change in the context and corresponding process changes. Examples of dynamic processes are processes for emergency management (a case study is presented in Section 1.3) and military forces deployment plans.

Some systems supporting structural changes of processes at run-time exist and are well supported by the research community [103, 48, 87, 104, 109, 139], but they do not automate the adaptation; a manual intervention of a domain expert is always required for adapting a faulty process instance at run-time. However, dynamic processes demand a more agile approach recognizing the fact that in dynamic environments process models quickly become outdated and hence require closer interweaving of modeling and execution [105]. To this end, our approach, named the SmartPM approach, will allow to *adapt automatically dynamic processes at run-time* when *unanticipated exceptions* occur, without the need to the define any recovery policy at design-time.

In the following section we show which classes of business processes currently exist on the market, we classify them on the basis of their degree of structure and we provide a systematic view of the different approaches and methodologies that have emerged to support different classes of processes. Moreover, we make clear the main requirements for modeling, executing and adapting dynamic processes.

1.2 The Spectrum of Process Management and Modeling Paradigms

When realizing a PMS based on executable process models, there is a variety of processes showing different characteristics and needs. On one hand, there exist well-structured and highly repetitive processes whose behavior can be fully pre-specified; on the other hand, many processes are knowledge-intensive and highly dynamic: typically, they can not be fully pre-specified and require loosely specified models [105]. In this section, we will analyze how operational and business processes can be classified on the basis of the degree of structuring they exhibit [63], which directly influences the level of automation, control, support and flexibility that they can provide.

1.2.1 Structured Processes

Figure 1.1(b) shows how processes can be classified on the basis of their “degree of structure” [63]. Traditional PMSs perform well with *structured processes* and controlled interactions between participants. Structured processes are characterized by a well defined structure in terms of activities to be executed and relations among them. They reflect highly repeatable routine work with low flexibility requirements (such as back-office financial transactions, manufacturing, production and administrative processes) that can be easily standardized and automated to increase efficiency [73]. A major assumption is that such processes, after having been modeled, can be repeatedly instantiated and executed in a predictable and controlled manner. All possible options and decisions (alternative paths) that can be made during process enactment are statically pre-defined at design time.

Structured processes with ad hoc exceptions have similar characteristics, but events and exceptions can occur that make the structure of the process less rigid and require process adaptation strategies. In the presence of expected exceptions, possible events and deviations that can be encountered are predictable and defined in advance, along with the specific handling logic, whereas the handling of unanticipated exceptions typically require (manual) structural process changes at run-time [104].

Structured processes can be completely captured by procedural process models that explicitly define the tasks and their execution constraints, participants, roles and input/output data. Most of classical process management environments deal with structured processes and are driven by imperative languages and procedural models, such as XPD L [146], WS-BPEL [92], Event-driven Process Chains (EPCs) [131], BPMN [143], UML Activity Diagrams (UML) [35] and YAWL nets [124]. All these languages mainly focus on the control-flow perspective and are widely used in research prototypes (e.g., YAWL [124]) and in open-source/free (e.g., jBPM [22], Apache ODE) and commercial products (e.g., Tibco Staffware Process Suite [125], Oracle BPEL Process Manager, IBM Process Manager [60]). While YAWL was directly defined starting from the Petri Nets formalism [88, 98], most of the aforementioned languages have been mapped to (variants of) Petri Nets (e.g., in [122, 34]) in order to provide a formal semantics and enable different verification techniques.

1.2.2 Loosely Structured Processes

In many application domains, pre-specifying the entire process model is not possible. A wide range of processes exhibit a loosely structured or semi-structured behavior (cf. unstructured processes with pre-defined fragments in Figure 1.1(b)), for which it is important to balance between flexibility and support (such as clinical guidelines and medical treatment procedures) and events and exceptions can occur that make the structure for the process significantly less rigid. While parts of the process logic are known at design-time, other parts are undefined or uncertain and can only be specified at run-time. For loosely specified processes, decisions regarding the specification of (parts of) the process have to be deferred to run-time.

Looseness is a characteristic of *Knowledge-intensive Processes*, that are processes characterized by being *non repeatable* (the models of two process instances may differ one another), *non predictable* (the course of the actions depends on context-specific parameters, whose values are not known a priori and may change during process execution) and *emergent* (the course of the actions only emerges during process execution, when more information is available) [105]. To some extent, knowledge-intensive processes and the looseness of their execution can be supported exploiting constraint-based process models. However, similarly to structured processes, loosely specified processes are still *activity-centric*, i.e., they focus on a set of activities that may be performed during process execution. Procedural models are not able to provide the degree of flexibility required in these settings as they may unnecessarily limit possible execution behaviors, with either over-specified or over-constrained models [86]. Process models should define tasks and their relationships in a less rigid manner, so that activities can be executed in multiple orders (or even multiple times) until the intended goals are achieved. Different modeling languages and management systems have been proposed and developed to meet these requirements, such as ADEPT2 [104], Flower [6], as well as service-oriented architectural solutions for the integration of different modeling approaches and PMS technologies [129].

Recent and ongoing work shows that declarative languages and models can be effectively used to increase the degree of flexibility for process specifications, still allowing to provide a good level of support. Languages such as ConDec [97] and DecSerFlow [134], supported by the Declare tool [96], propose a declarative constraint-based approach for modeling, enacting and monitoring business processes. Instead of strictly and rigidly defining the control-flow of process tasks using a procedural language, they exploit the concept of control-flow *constraints*, defined as Linear Temporal Logic (LTL) formulae [137], for the specification of relationships among tasks (generally classified as existence, choice, relation, negation and branching constraints over process activities). Constraints implicitly define possible execution alternatives by prohibiting undesired execution behavior, and they reflect policies and business rules to be satisfied and followed in order to successfully perform a process and achieve the intended goals.

1.2.3 Unstructured Processes

Unstructured processes are characterized by a low level of structuring and an high degree of flexibility. Process participants decide on the activities to be executed

as well as their execution order, and the structure of a process thus dynamically evolves and strongly depends on user decisions made during process execution. These processes reflect knowledge work and collaboration activities driven by rules and events, for which no predefined models can be specified and little automation can be provided. Knowledge workers rely on their experience and capabilities to perform ad hoc tasks on a per-case basis and handle unexpected events and changes in the operational context.

For processes with these characteristics only their goal is known a priori and they can not be fully pre-specified at a fine-grained level at design time. In practice, approaches that focus on the role of *data* as main driver for process execution and activity coordination are required. In this direction, Adaptive Case Management (ACM) [147] has emerged as a way for supporting unstructured, unpredictable and unrepeatable business cases. ACM adopts a *data-centric* (rather than an *activity-centric*) approach, focusing on the concept of *case* (an insurance claim, a customer purchase request, patient medical records, etc.) as primary object of interest, and the progress of the case itself is driven by the availability, values, changes and evolution of data objects and their dependencies. Each execution of a case management process involves a particular situation (the case) and a desired outcome (or goal) for that case, and the determination of actions to take in each case involves the exercise of human judgement and decision-making at run-time. In 2009, the Object Management Group (OMG) issued a Request for Proposal [94] for a meta-model extension to BPMN 2.0 to support modeling of case management processes but, to date, there exist no standards for supporting case management process modeling and execution.

Another approach for dealing with unstructured processes is the *object-aware approach* proposed in [67]. For object-aware processes, the information perspective is predominant and captures object types, their attributes, and their interrelations, which together form a data structure or *information model*. In accordance to data modeling, the modeling and execution of processes can be based on two levels of granularity: object behavior and object interactions. At run-time, the different object types comprise a varying number of inter-related object instances, that evolve according to their specified behavior and interaction models. Process execution and possible actions are thus related to objects and their states [66]: the enabling of a process step or action does not directly depend on the completion of preceding steps (i.e., on the control-flow as for activity-centric approaches), but rather on the changes and evolution of object states and relations.

Similarly, artifact-centric models [91, 56] aim at providing a declarative data-centric modeling approach where, rather than prescribing control-flow constraints between tasks (as in process-centric models), process specification and enactment are driven by data dependencies and evolutions of business entities.

1.2.4 Dynamic Processes

Both structured and loosely specified processes are activity-centric; i.e., they are based on a set of activities that may be performed during process execution. However, the class of *dynamic processes* is transversal with respect to the classification proposed in [63]. These processes represent activities in highly dynamic situations and

unforeseen exceptions (e.g., emergency management scenarios) and that are executed in a world with little structure and possibly imperfect information. The scenario dictates who should be involved and who is the right person to execute a particular step of the process, and collaborative interactions among the users typically is a major part of such processes.

When trying to support and implement dynamic processes through a PMS, a complete integration of processes, data and users has to be achieved. Usually, the structure of a dynamic process can be completely captured with a *procedural process model* that explicitly defines the tasks and their execution constraints. However, a dynamic process is thought to be enacted in pervasive and highly dynamic scenarios, where exceptions and exogenous events “are not the exception but the rule”. Therefore, there is the need to explicitly represent the *contextual data* describing the scenario in which the process will be enacted, *constraints* that evaluate if a particular activity can be applied in a specific state of the contextual scenario and *dependencies* to specify execution dependencies between activities.

The execution of process tasks may affect the data associated to the contextual scenario, and it may happen that the *undesirable outcome of some activities* causes changes in the contextual environment that may prevent the achievement of the business goals. The same risk comes from *exogenous events*, that represent external events that come from the environment and correspond to changes in the data associated to the contextual scenario. In both cases, in order to make a dynamic process adaptable to the contextual environment, a PMS implementation should automatically (i) *detect exceptional situations*, (ii) *derive at run-time* and (iii) *correctly apply* the recovery procedures necessary to handle them, by involving some form of reasoning on the available process tasks and contextual data. In particular, the third point imposes a strong requirement: *any structural change made to the process instance at run-time must not violate the process model correctness and process instance execution*.

In order to meet these requirements, in Chapters 3 and 4 we propose a modeling language, an approach and a PMS realization named SmartPM that allow to perform automatic adaptation of those processes at run-time. In the meanwhile, an extensive case study involving a real dynamic process is shown in the following section.

1.3 Case Study

The SmartPM system has been validated in laboratory tests through the use of a real process of the Italian Railway company (that is “Reti Ferroviarie Italiane”). Specifically, let us consider the emergency management scenario described in Fig. 1.2(b). It concerns a train derailment and depicts a map of the area (as a 4x4 grid of locations) where the disaster happened. For the sake of simplicity, we suppose that the train is composed of a locomotive (located in *loc33*) and two coaches (located in *loc32* and *loc31* respectively). The goal of an incident response plan defined for such a context is to evacuate people from the coaches located in *loc32* and *loc31* and to take pictures for evaluating possible damages to the locomotive, located in *loc33*.

Thus, a response team can be sent to the derailment scene. The team is composed of four first responders (in the remainder, we refer to them as *actors*) and two robots,




Figure 1.2. Area (and context) of the intervention.

initially located in *loc00*. We assume that actors are equipped with mobile devices (for picking up and executing tasks) and provide specific capabilities. For example, actor *act1* is able to extinguish fire and take pictures, while *act2* and *act3* can evacuate people from train coaches. The two robots, instead, may remove debris from specific locations. Each robot has a battery and each action consumes a given amount of battery charge. When the battery of a robot is discharged, actor *act4* can charge it. Moreover, the battery charge consumption amounts are provided for each action (e.g., taking pictures in a given location consumes 1 unit of battery charge, removing debris consumes 3 units, etc.).

In order to carry on the overall process, all the actors/robots need to be continually inter-connected. The connection between mobile devices is supported by a network provided by a fixed antenna (whose range is limited to the dotted squares in Fig. 1.2(b)), and the robots *rb1* and *rb2* can act as wireless routers for extending the network range in the area. A robot provides a connection limited to the locations adjacent (in any direction) to its position. Each robot can move in the area, but it is constrained to be always connected to the main network. This is guaranteed if the intersection between the squares covered by the main network and the squares covered by the robot connection is not empty. A robot connected to the main network can act as a “bridge”, allowing the other robot to be connected through it to the main network.

Collected information is used for defining and configuring at run-time an incident response plan, defined by a contextually and dynamically selected set of activities to be executed on the field by first responders. A possible concrete realization of the incident response plan is shown in Figure 1.2(a). The process is composed by three parallel branches with tasks that instruct first responders to act for evacuating people from train coaches, to take pictures and to assess the gravity of the accident. Despite the simple structure of the incident response plan, the high dynamism of the operating environment can lead to a wide range of exceptions. In general, for dynamic processes there is not a clear, anticipated correlation between a change in

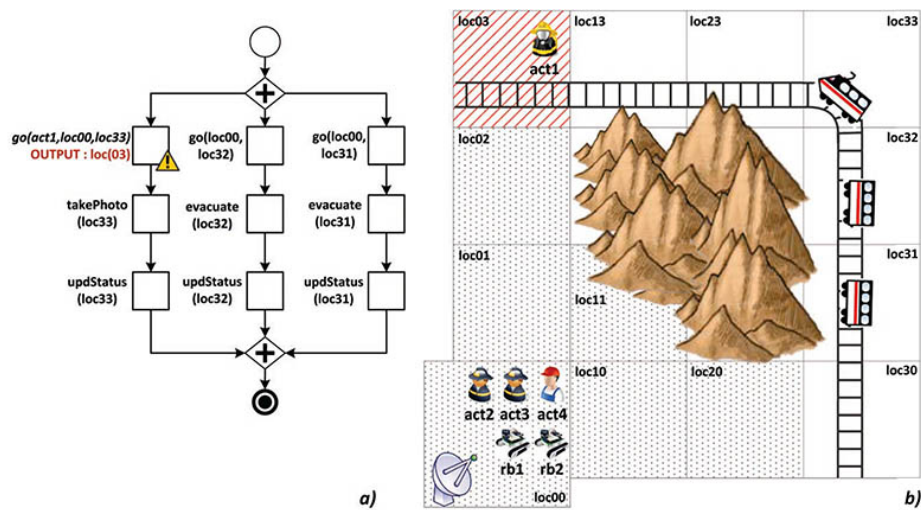


Figure 1.3. Effects of a task with an outcome different from the one expected.

the context and corresponding process changes. Suppose, for example, that the task $go(loc00, loc33)$ is assigned to actor $act1$ (cf. Fig. 1.3(a)), which reaches instead the location $loc03$ (cf. Fig. 1.3(b)). This means that $act1$ is now located in a different position than the desired one, and s/he is out of the optimal network range. Since all the actors/robots need to be continually inter-connected to execute the process, the PMS has to find a recovery procedure that first instructs the robots to move in specific positions for maintaining the network connection, and then re-assign the task $go(loc03, loc33)$ to $act1$.

Even this very simple example shows that the same failure may require significantly different adaptation activities depending on the current context. For example, if robots $rb1$ and $rb2$ have enough battery charge, the PMS can instruct first $rb1$ to move in $loc03$ (cf. Fig. 1.4(a) and Fig. 1.4(b)), in order to re-establish the connection of actor $act1$ with the network. Then, robot $rb2$ can be instructed to reach $loc23$ (cf. Fig. 1.4(a) and Fig. 1.4(c)) and to broaden the network range and make $loc33$ (the expected destination of $act1$) covered by the network connection. Finally, the task $go(loc03, loc33)$ can be effectively reassigned to $act1$ (cf. Fig. 1.4(a) and Fig. 1.4(d)).

After having executed the recovery procedure, the enactment of the main process can be resumed to its normal flow. The execution of a dynamic process can be also prevented by the occurrence of external events (a.k.a. exogenous events). These are events coming from the external environment that change asynchronously some contextual properties of the scenario in which the process is under execution. For example, in Fig. 1.5(a) we are supposing that a rock slide has collapsed on location $loc31$. The presence of a rock slide modifies the contextual properties of the scenario in a way not expected when the dynamic process was designed, by possibly jeopardizing the correctness of the dynamic process itself. Since a dynamic process - by definition - has to be *adaptable to the context*, the PMS needs to find a recovery procedure that allows to remove the rock slide from $loc31$ by maintaining all the process participants inter-connected. A possible solution is shown in Fig. 1.5(b) and Fig. 1.5(c), and consists in moving the robot $rb1$ in $loc31$ and letting it remove




Figure 1.4. Synthesis and enactment of the recovery process after an undesirable outcome.

debris.

It is unrealistic to assume that the process designer can pre-define all possible compensation activities for dealing with these exceptions (apparently simple), since the process may be different every time it runs and the recovery procedure strictly depends on the actual contextual information (the positions of actors/robots, the range of the main network, the battery level of each robot, etc.). In the worst case, the number of recovery processes to pre-define may depend to all the possible combinations of contextual information. For the same reason, it is also difficult to manually define an ad-hoc recovery procedure at run-time, as the correctness of the process execution is highly constrained by the values (or combination of values) of contextual data.

The main purpose of the SmartPM system is to develop a PMS providing automatic mechanisms that, starting from a process model, are able to adapt the process instance without explicitly defining handlers/policies to recover from exceptions and exogenous events and without the intervention of domain experts.

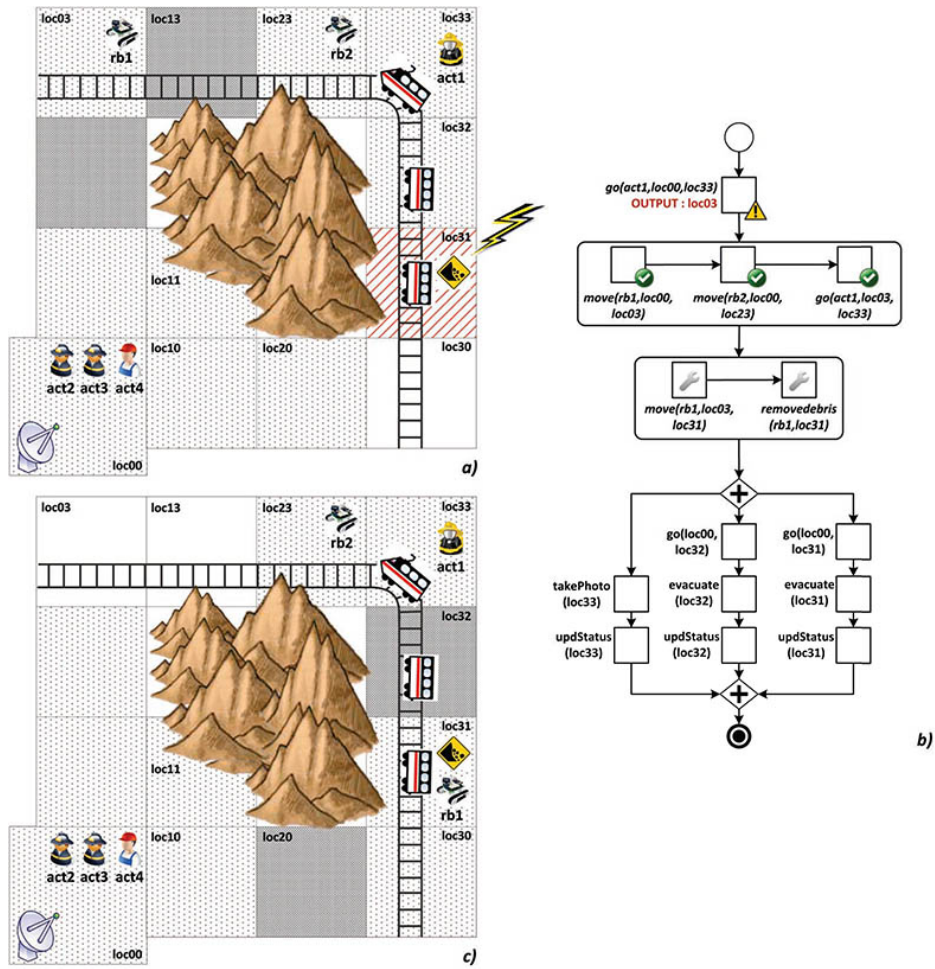


Figure 1.5. Synthesis and enactment of the recovery process after the occurrence of an exogenous event.

Chapter 2

State of the Art

Process adaptation refers to the ability of a PMS to modify its behavior according to environmental changes and exceptions that may occur during process execution. If not detected and handled effectively, exceptions can result in severe impacts on the cost and schedule performance of PMSs [64].

Nowadays, PMSs provide wide support for different modeling styles and for all phases of the process life-cycle, from the specification and enactment to the verification, monitoring and analysis of process models [36]. In addition, PMSs provide tools for modelling business processes that are predictable and repetitive. However, in many real world scenarios, enabling process adaptation is crucial for any PMS, as implemented processes may have to be adapted to deal with changing environments and evolving needs.

In this chapter, we focus on the *exception handling* perspective provided by current state-of-the-art PMSs, by first showing some of the best known techniques for exception handling and then by reviewing some of the most interesting current commercial and academic prototype PMSs in relation to their approaches to process adaptation. Finally, in Section 2.2 we analyze a number of techniques from the field of Artificial Intelligence (AI) that were applied to BPM, with the purpose to facilitate automatic adaptation of a business process at run-time.

2.1 Process Adaptation

Over the last years, there was a trend in providing PMSs with a growing support for adapting business processes. Proposed approaches can be analyzed considering to what extent users are involved in the process of defining exception conditions and handling policies, and the degree of automation provided in the exception resolution and process adaptation stages.

This section is thought for discussing the current state-of-the-art in process adaptation and exception handling (cf. Section 2.1.1) and for reviewing some of the most interesting current commercial and academic prototype PMSs in relation to their approaches to process adaptation (cf. 2.1.2). In Section 2.1.3 we provide a comparative analysis between the existing techniques and the **SmartPM** approach we are going to present in this thesis.

2.1.1 Exception Handling Techniques

One of the first attempts to define and categorize exceptions in PMSs was proposed in the nineties in [39], where possible exceptions were classified as *basic failures*, *application failures*, *expected exceptions* and *unanticipated exceptions*. While basic failures (which reflect a failure at system level, e.g., DBMS or network failure) and application failures (corresponding to the failure of an application implementing a given task) are generally handled at the system and application levels, PMSs are in charge to provide support for *exception handling*.

An ***expected exception*** is an exception that can be planned at design-time, i.e., a process designer can provide an exception handler which is invoked during run-time to cope with the exception itself. In [13], Casati identifies 4 potential sources for expected exceptions:

- *Process Exceptions*. They can occur during the process enactment, and are strictly related to *task failures*. More in detail, a task can fail due to an abnormal termination of the invoked application or web service implementing the specific task, or because of a negative termination of the task itself (which returns an output different from the expected one).
- *Constraint Violations*. They refer to violations of constraints over data (e.g., data required for task execution might be missing), tasks (e.g., in terms of pre/post-condition of a task not satisfied before/after task execution) and resources (e.g., unavailability of resources during process execution).
- *Temporal Exceptions*. They can be associated with deadlines, and upon deadline expiration an exception is launched.
- *External Exceptions*. They happen when an external event may affect the control/data flow of the process under execution.

To enable an expected exception to be detected at run-time, a PMS has to notice its occurrence during process enactment. To this end, a process designer needs to associate pre-specified process models with exception handlers at design-time. Along this line, in [71] the authors propose a patterns-based approach to exception systematization, defining a general classification framework and language for exception handling. Specifically, exception handlers are modeled as alternative branches of the process model that add or cancel behavior to the normal flow of the process instance under execution.

Among early attempts to address the problem of exception handling, it is worth mentioning the transaction oriented Workflow Activity Model WAMO [38]. The model allows the designer to define expected exceptions which may arise during process execution and to derive compensation tasks for exception handling. The specification of compensation activities avoids the designer to explicitly model all possible process execution alternatives in presence of exceptions, and allows the system to automatically control the reliable execution of exceptions and failures. Handling capabilities depend on task properties and range from simple re-execution of failed activities to the execution of complex compensation strategies.

In [51] the authors proposed Self-Adaptive Recovery Nets (SARNs), an extended Petri net model for specifying exceptional behavior in PMSs at design time. SARNs can adapt the structure of the underlying Petri net model at run time to handle

pre-specified exceptions. Process adaptation is achieved by applying high-level recovery policies (such as skipping, redoing or compensating tasks when specific events occur) defined at design time for single tasks or sets of tasks.

In [11], authors argue that the majority of exceptions are asynchronous (they can occur at any time during the process execution, rather than before/after a task execution), and that they can not be suitably represented within the process model. Therefore, they advocate that the only way to handle exceptions is programmatically via Event-Condition-Action (ECA) rules. ECA rules have the form “on *event* if *condition* do *action*” and specify to execute the *action* (i.e., the exception handler) automatically when the *event* happens (i.e., when the exception is caught), provided the a specific *condition* holds.

ECA rules represent a good way for separating the graphical representation of the process with the “exception handling flow”. A similar principle has been applied in YAWL [124]. In YAWL, for each exception that can be anticipated, it is possible to define an exception handling process, named *exlet*, which includes a number of exception handling primitives (for removing, suspending, continuing, completing, failing and restarting a workitem/case) and one or more compensatory processes in the form of *worklets* (i.e., self-contained YAWL specifications executed as a replacement for a workitem or as compensatory processes). Exlets are linked to specifications by defining specific rules (through the Rules Editor graphical tool), in the shape of *Ripple Down Rules* specified as “if condition then conclusion”, where the condition defines the exception triggering condition and the conclusion defines the exlet.

The above approaches use to deal with exceptions by attempting to anticipate all possible failures. This greatly complicates the process models and thereby obscures the main “preferred” process if the number of exceptions to define is too large [64]. Moreover, in many cases, exceptions can only be modeled at the technical level, although their handling requires organizational knowledge and thus should be expressed in semantic process models as well. Hence, authorized users should be allowed to situationally adapt single process instances during run-time to cope with unanticipated exceptions [105].

An ***unanticipated exception*** can be detected during the execution of a process instance, when a mismatch between the computerized version of the process and the corresponding real-world business process occurs. To cope with those exceptions, a PMS is required to allow structural adaptation of its corresponding process model at run-time.

In [105] is presented a basic taxonomy and respective patterns for applying structural changes to running process instances in case of unanticipated exceptions. Specifically, a structural change can be applied on the *state of the process instance* or/and on the *structure of its process model* and concerns the application of *change primitives* to the process model (i.e., by adding or deleting single tasks or transitions) or of *high-level change operations* (i.e., by inserting, deleting or moving activities or entire process fragments). Another dimension concerns the *degree of automation* in applying a change. The majority of the approaches dealing with unanticipated exceptions provide a *manual* definition of a structural change [119, 141, 4, 108, 104]; an expert user suspends the process instance, loads it into a process editor and manually adjusts its structure.

Since this procedure is quite complex for non expert end-users, a growing number of systems [109, 139, 84] is providing a *semi-automatic* way for assisting end-users in defining structural changes at run-time by exploiting the *Case-Based Reasoning* (CBR) approach [68]. CBR is a paradigm based on reasoning and machine learning, and solves an emerging problem by drawing on past experiences and by adapting respective solutions to the new problem situation. Problems and their solutions are described as *cases* and stored in *case-bases*. Each problem that has been solved provides a new case that is stored for future reuse, making learning a natural side-effect of the reasoning process. When using CBR, the end-user is required to provide a complete problem specification at design-time for case retrieval.

By the way, even systems that do support manual and semi-automatic dynamic process model modification at run-time do not help determine the best response to a given exception, leaving this often difficult decision to end-users [64]. In order to relieve end-users from building complex adaptation tasks, there are few systems [83, 74, 40, 43] that provide approaches for *automatically* adapting a process in presence of unanticipated exceptions. While we will present the goal-based approaches [74, 40, 43] in Section 2.2, it is interesting to say a few words about the work [83]. It concerns a recent framework dealing with automatic adaptation of processes through CBR. The method employs process adaptation cases that record adaptation episodes from the past. The recorded changes can be automatically transferred to a new process that is in a similar situation of change. The case-based adaptation method includes the so-called *anchor mapping algorithm* which identifies the parts of the target process where to automatically apply the changes.

2.1.2 Analysis of Existing PMSs

Generally, existing commercial PMSs provide only basic support for handling exceptions in a proprietary manner [148]. They typically require the process model to be fully defined before it can be instantiated, and changes must be incorporated by modifying the model statically:

- *COSA* [20] allows to associate sub-processes to external “triggers” and events. All events, sub-processes and adaptation policies must be defined at design-time, although models can be modified at run-time (but only for future instantiations). If a task fails, the associated activity can be rolled back or restarted. When a deadline expires, a compensating activity is triggered. *COSA* also allows manual ad-hoc run-time adaptation by using change patterns [138] such as reordering, skipping, repeating, postponing or terminating steps.
- *Tibco iProcess Suite* [125] provides constructs called “event nodes”. For each event node, a separate pre-defined exception handler can be activated when an exception occurs at that point. If no handler exists for the specific exception, the identified exception is forwarded into a “default exception queue” where it may be manually handled. Moreover, the designer can decide to manually skip tasks at run-time.
- *WebSphere MQ Workflow* [60] supports deadlines and, when they expire, some pre-defined exception handlers are invoked and/or notification messages are sent to the process administrator, which can manually suspend, restart or terminate processes, as well as s/he can reallocate tasks.

- *SAP Workflow* [85] allows to provide exception events for checking tasks pre- and post-constraints and for waiting until an external trigger occurs. Exception handlers are pre-defined at design-time and, when an exception occurs at run-time, the process administrator can manually select the most appropriate recovery procedure from the pre-specified ones. When an exception handler is found, all tasks in the block where the exception is caught are cancelled.

The main issue with the above approaches is that they provide solid support for transaction-level exceptions and very few support for exceptions at process or instance level. Moreover, exceptions are mainly thrown when some deadline expires, and also the behavior of those systems after a deadline is occurred depends mainly by a process administrator [2]. Research in process adaptation has born mainly for overcoming the above limitations, with the purpose to increase the degree of automatic adaption of business processes. There have been a number of academic prototypes developed in the last decade, and we limit the discussion to the more popular and recent systems and prototypes¹:

- The *OPERA* prototype [17] borrows its main ideas from a combination of programming language concepts and transaction processing techniques, adapting them to the special characteristics of PMSs. OPERA was one of the first systems that integrate language primitives for exception handling into PMSs. It has a modular structure in which activities are nested. When a task fails, its execution is stopped and a pre-specified exception handler is launched for dealing with that kind of exception. If the handler cannot solve the problem, it propagates the exception up the activity tree; if no handler is found the entire process instance aborts.
- In the *ADOME* system [18], the problem of exception handling is addressed adopting an integrated event-driven approach that covers exception detection, exception resolution, reuse of exception handlers, and automated resolution of expected exceptions. As exceptions are identified by specific events, Event-Condition-Action (ECA) rules are used to dynamically bind exception handlers to exception conditions at design time, and enable automatic exception detection and resolution at run-time.
- An interesting approach for dealing with process adaptation is provided by the *YAWL* system [124]. Exception handling capabilities provided by YAWL were designed and implemented starting from the conceptual framework for workflow exception handling presented in [113]. For each exception that can be anticipated, it is possible to define at design time exception handling processes (*exlets*) linked to process specifications by defining triggering conditions and rules in the form of Ripple Down Rules [1]. Handling processes, which include exception handling primitives (for removing, suspending, continuing, completing, failing and restarting a work item/case) and one or more compensatory processes, are dynamically selected and incorporated at run-time in process instances. More details about YAWL are given in Chapter 7.
- Strong support for adaptive process management and exception handling is provided by the *ADEPT* system. *ADEPT_{flex}* offers modeling capabilities to explicitly define pre-specified exceptions, and supported changes of process

¹A more detailed analysis can be found in [2].

instances to enable different kinds of ad-hoc deviations from the pre-modeled process schemas in order to deal with run-time exceptions [103]. These features have been extended and improved in *ADEPT2* [108, 104], which provides full support for the structural process change patterns defined in [138], and in *ProCycle*, which combines ADEPT2 with conversational case-based reasoning (CCBR) methodologies [109]. CCBR is an extension of CBR, where the user is not required to provide a complete a priori specification for case retrieval. Specifically, process instance changes are manually performed and stored as *cases* in a case-base, to be retrieved and reused to perform similar changes in the future. Users are supported in finding relevant cases and adapting processes by taking into account how similar events were previously handled and reusing information about similar changes applied to previously executed instances.

- The *AdaptFlow* prototype [48] supports ECA rules-based detection of exceptions and the dynamic adaptation of process instances, although each adaptation must be confirmed manually by a process administrator before it is applied. If the exception handler does not satisfy the process administrator, s/he can manually deal with the exception. The prototype has been designed as an overlay to the ADEPT system, providing dynamic extensions.
- Similarly, *AgentWork* [87] relies on ADEPT and exploits a temporal ECA rule model to automatically detect logical failures and enable both reactive and predictive process adaptation of control- and data-flow elements. Here, exception handling is limited to single tasks failures, and the possibility exists for conflicting rules to generate incompatible actions, which requires manual intervention and resolution.
- Conversational CBR has been also applied in the *CBRFlow* system [139] for supporting adaptation of pre-defined process models to changing circumstances by allowing (manual) annotation of business rules during run-time via incremental evaluation by the user. Thus users must be actively involved in the inference process during each case.

2.1.3 Discussion

All the approaches analyzed so far share a common background. At design-time the process modeler identifies possible exceptions that can occur, explicitly defines exception conditions, triggering events and handling policies, and integrates such information into the the process model. At run-time exceptions are detected and handled either manually or (semi-)automatically, adapting affected processes at the instance level.

A good way for understanding and comparing the kind of adaptation we are proposing with the other existing approaches is to consider the `try-catch` approach used in some programming languages such as Java². If the `try` block encloses the code that might throw an exception, the `catch` is the definition of the possible exceptions and the specification, defined at design-time by the process modeler, of how to deal with them.

²<http://www.java.com/en/>

For expected exceptions, both the `try` and the `catch` blocks need to be completely pre-specified, i.e., for each kind of exception that is envisioned to occur, a specific exception handler is defined. Systems like [20, 125, 60, 85, 17, 18, 124] provide mainly this kind of support to process adaptation.

On the contrary, for unanticipated exceptions, the `try` is automatically derived as the situation in which the PMS does not adequately reflect the real-world process anymore. As a consequence, one or several process instances have to be adapted, and the `catch` block should include those recovery procedures required for realigning the computerized processes with the real-world ones. The common strategy used by the current adaptive PMSs [103, 48, 87, 104, 109, 139] is to *manually* or *semi-automatically* define at run-time the `catch` block, by devising an ad-hoc recovery procedure that deals with the specific exception captured.

Our proposal - instead - aims at *synthesizing at run-time* the `catch` block, without the need of any manual intervention at run-time. This because the design-time specification of all possible compensation actions requires an extensive manual effort for the process designer, which has to anticipate all potential problems along with possible ways to overcome them. This is particularly true in real-world and dynamic scenarios, where the process designer often lacks the needed knowledge to model all the possible contingencies at design-time, or this knowledge can become obsolete as process instances are executed and evolve, by making useless his/her initial effort. For the same reason, analyzing and defining these adaptations “manually” becomes time-demanding and error-prone in real-world scenarios. Indeed, the designer should have a global vision of the application and its context to define appropriate recovery actions, which becomes complicated when the number of relevant context features and their interleaving increases.

2.2 AI-based Process Adaptation

The Artificial Intelligence (AI) community has been involved with research on process management for several decades [89]. While BPM has concentrated on business and manufacturing processes, the AI community has been motivated primarily by domains that involve active control of computational entities and physical devices (e.g., robots, antennas, satellites, computer networks, software agents, etc.). Despite their differing origins and emphases, there is much overlap between the objectives, requirements, and approaches for process management within these two communities. Specifically, AI technologies can play an important role in the construction of the PMS engines that manage complex processes, while remaining robust, reactive, and adaptive in the face of both environmental and tasking changes. An interesting report that describes how techniques from the AI community could be leveraged to provide several of the advanced process management capabilities envisioned by the BPM community is shown in [89].

The advantage of integrating AI planning techniques [90] for several applications in the field of BPM has long been acknowledged. For example, in [61] the authors take a broad view of the problem of adaptive PMSs, and show that there is a strong mapping between the requirements of such systems and the capabilities offered by AI techniques. The work describes how planning can be interleaved with

process execution and plan refinement, and investigates plan patching and plan repair as means to enhance flexibility and responsiveness. The approach presented in [100] highlights the improvements that a legacy workflow application can gain by incorporating planning techniques into its day-to-day operation. The use of contingency planning [32] to deal with uncertainty increases system flexibility, but it does suffer from a number of problems. Specifically, contingency planning is often highly time-consuming and does not guarantee a correct execution under all possible circumstances.

Also the use of planning techniques for process adaptation purposes is not strictly new. One of the first work dealing with this research challenge is [5]. It discusses the use of an intelligent assistant based on planning techniques, whose purpose is to suggest compensation procedures or the re-execution of activities if some failure arises during the process execution. The intelligent assistant makes use of a meta-level knowledge incorporated into the process model.

A number of *goal-based approaches* exist for enabling automatic process instance change in case of emerging exceptions. Those approaches [74, 40, 43] explicitly formalize an output for the process to be executed in terms of a process *goal*. Starting from the process goal, they are able to automatically derive a process model (i.e., the activities to be performed and their execution order) that complies with the goal itself. In addition, if an activity failure occurs at run-time and leads to a goal violation, the process instance is adapted accordingly. Specifically:

- EPOS [74] presents a general, incremental replanning algorithm, which allows to automatically adapting process instances when process goals change. The adaptation consists in synthesizing a new process specification that reflects the changes in the process goal.
- A similar approach is proposed in [43], where the authors present a concept for dynamic and automated workflow re-planning that allows recovering from task failures. To handle the situation of a partially executed process, a multi-step procedure is proposed that includes the termination of failed activities, the sound suspension of the process, the generation of a new complete process definition and the adequate process resumption.
- A new life cycle for workflow management based on the continuous interplay between learning and planning is proposed in [40]. The approach is based on learning business activities as planning operators and feeding them to a planner that generates a candidate process model that is able of achieving some business goals. If an activity fails during the process execution at run-time, an alternative candidate plan is provided on the same business goals.

Planning techniques are also used in [42] to define a self-healing approach for handling exceptions in service-based processes and repairing faulty activities. The purpose of [42] is to provide a tool that analyzes the repairability of service-based processes and generates repair plans based on a set of design-time repair actions. At design time, the process description is augmented with information about: (i) the data dependencies within the process, (ii) the desired process results, and (iii) available repair actions for each process activity. In addition, an heuristic-based approach is used for reasoning about the repairability of the process schema, by checking if it includes some design flaw that may possibly affects the repairability

of its activities. When a failure is detected during the process execution, a *diagnosis tool* determines what caused the failure and where (i.e., in which point of the process) the error took place. Such knowledge, obtained by the diagnosis tool or - if needed - with the support of a human process modeler, is used together with the repairability information inferred at design time for generating a repair plan dealing with the exception. Repair is specified as a planning problem whose goal is to build a plan consisting of recovery actions that, after being executed, should recover the faulty process instance. The repair plan is generated by taking into account constraints posed by the process structure and by applying or deleting repair actions taken from a given *generic repair plan*. The generic repair plan is built by considering all possible applications of the recovery actions in a specific state.

The work [9] proposes a goal-driven approach to business process adaptation in service-based applications. The authors provide an approach that allows to automatically adapt business processes to run-time context changes that impede achievement of a business goal. The execution of the business process and the evolution of its context are continuously monitored against the desired model of the application, which is based on the notion of goal and on service annotations describing how services contribute to goal achievement. The model encodes also the business policies defined over the elements of the domain; each contextual property is modeled with a state transition system capturing all possible values of the specific property, and their evolution after a precondition violation or the occurrence of an exogenous event. The adaptation mechanism is based on service composition via automated planning techniques [90]: whenever a deviation is detected (that results in a violation of some business policy), a specific planner will be in charge to generate a composition of available services that achieves the adaptation goals in compliance with the business policies. In this approach, compensation activities that compose a recovery procedure are not required to be explicitly represented at design-time, but are automatically derived from the various aspects of the environment (the contextual properties, the business policies, the available services, etc.).

Another interesting work dealing with process interference is [126]. Process interference is a situation that happens when several concurrent business processes that depend on some common data are executed in a highly distributed environment. During the processes execution, it may happen that some of these data are modified causing unexpected or wrong business outcomes. To overcome this limitation, the work [126] proposes a run-time mechanism which uses *Dependency Scopes* - for identifying critical parts of the processes whose correct execution depends on some shared variables - and *Intervention Processes* for solving the potential inconsistencies generated from the interference. In [127], the authors present an algorithm for automating the discovery of critical sections and the generation of dependency scopes, whereas intervention processes need to be manually defined at design time. One of the future work claimed by the authors concerns to devise a technique for the automatic synthesis of intervention processes through a domain independent planner based on CSP techniques. Details about the planner and the algorithm used for synthesizing a plan are shown in [62].

2.2.1 Discussion

If compared with the above works, our approach provides some unique features in dealing with unanticipated exceptions for dynamic processes:

- The goal-based approaches [74, 40, 43] apply planning techniques to automatically derive a process model from business goals, and to repair it during run-time if goals change by replanning the process specification. However, the structure of a dynamic process usually derives from a pre-defined high-level procedure, and it is difficult to extract a goal representing the correct execution of the process. Moreover, current planning methods do not still cover adequately all relevant modeling aspects (e.g., treatment of loops, appropriate handling of data flow, etc.). But the major issue lies in the replanning stage used for adapting a faulty process instance. In fact, dynamic processes usually involve real human participants acting in pervasive scenarios, and to re-define completely the process specification at run-time when the process goal changes (due to some activity failure) means to completely revolutionize the work-list of tasks assigned to the process participants. On the contrary, our approach adapts a running process instance by modifying only those parts of the process that need to be changed/adapted by keeping other parts stable.
- In the work [42], it is required to explicitly define repair activities at design-time. But, for a dynamic process, it is not an easy task to predict how the process instances will unfold and to understand which recovery actions effectively will be required for adapting the processes. Instead, in our approach, the recovery procedure will be built by composing tasks stored in a specific repository. The repository contains both tasks used for defining the specific process instance under execution and other tasks built on the same contextual scenario and possibly used in past executions of the process. This means we don't explicitly define at design time any recovery action, by letting an external planner to decide which tasks are really required for building the recovery procedure. By the way, the real bottleneck of the approach presented in [42] consists in the presence of a generic recovery plan from which an ad-hoc repair plan is derived when an exception occurs. The authors exploit the disjunctive logic programming system DLV [70] to search for successful repair plans starting from the generic one. However, given n repair actions, there exist, in the worst case, at least $O(n^n)$ possible repair plans, since all possible permutations of repair actions have to be considered. For avoiding such a complexity authors apply heuristics that limit the search space. The drawback is that some successful repair plans may be possibly not computed.
- The work [9] presents the interesting feature of providing services with non-deterministic outcomes. However, in this approach, service operations as well as context properties are modeled through state transition systems, by requiring a considerable manual effort at design-time. Furthermore, the approach does not support numeric variables and would suffer of state explosion and poor performance if monitored variables range over large domains. The main difference between the SmartPM approach and the one presented in [9] is that in SmartPM the recovery procedure is synthesized at runtime, without the need to define any recovery policy at design-time. On the contrary, in [9], the

process modeler is required to define some “business policies” for detecting the exceptions.

- The work presented in [126] is the closest to the SmartPM approach, but it does not still provide automatic mechanisms for the building of the recovery procedure (however, it has been claimed to be one of the main future direction of this work) and it only considers external events as a source of failure (we focus on both external events and tasks failures).

Chapter 3

The SmartPM Approach

Approaches enabling automated process instance adaptation at run-time aim at reducing error-prone and costly manual ad-hoc changes, and thus at relieving users from complex adaptations tasks [105]. As a prerequisite for such automated changes, a PMS must be able to:

- automatically detect exceptional situations;
- derive at run-time a recovery procedure to handle them;
- correctly apply the recovery procedure to the process instance involved in the specific exception.

To address this research issue, in the following two chapters we present SmartPM (that is the acronym for *Smart Process Management*), which is a model and a prototype PMS featuring a set of techniques providing support for automatic adaptation of processes at run-time. Such techniques are able to automatically adapt process instances without explicitly defining handlers/policies to recover from tasks failures and exogenous events at run-time and without the intervention of domain experts. While in Chapter 4 we will describe in detail the architecture and the technological aspects of the SmartPM software prototype, in this chapter we focus on the theoretical aspects of our general approach for adaptation of dynamic processes.

To accomplish this, we make use of well-established techniques and frameworks from Artificial Intelligence, such as *Situation Calculus*[107], *IndiGolog* [25] and *automatic planning* [90]. After presenting in Section 3.1 an overview of our general approach, in Section 3.2 we will discuss some preliminary notion around these challenging topics. Then, in Section 3.3 we show how we explicitly formalize processes in situation calculus and IndiGolog and in Section 3.4 we discuss how this formalization can be used for devising a technique that automatically detect and recover from failures. Finally, in Section 3.5 we present three different adaptation mechanisms developed within the SmartPM prototype.

3.1 Overview of the Approach

SmartPM adopts a *service-based* approach to process management, that is, tasks are executed by services (that could be software applications, human actors, robots). In SmartPM a process model is defined as a set of n task definitions, where each task t_i

can be considered as a single step that consumes input data and produces output data. Data are represented through a set F of *fluents* f_j whose definition depends strictly on the specific process domain of interest. In AI, a fluent is a condition that can change over time. Such fluents can be used to constrain the task assignment (in terms of *task preconditions*), to assess the outcome of a task (in terms of *task effects*) and as guards into the expressions at decision points (e.g., for cycles or conditional statements).

Choosing the fluents that are used to describe each activity falls into the general problem of *knowledge representation*. To this end, the environment, services and tasks are grounded in domain theories described in Situation Calculus [107]. Situation Calculus is specifically designed for representing dynamically changing worlds in which all changes are the result of the tasks' execution. Processes are represented as IndiGolog programs. IndiGolog [25] allows for the definition of programs with cycles, concurrency, conditional branching and interrupts that rely on program steps that are actions of some domain theory expressed in Situation Calculus. The dynamic world of SmartPM is modeled as progressing through a series of *situations*. Each situation is the result of various tasks being performed so far. Fluents may be thought of as “properties” of the world whose values may vary across situations.

SmartPM provides mechanisms for adapting process schemas that require no pre-defined handlers. To this end, a specialized version of the concept of adaptation from the field of agent-oriented programming [27] is used. Specifically, adaptation in SmartPM can be seen as reducing the gap between the *expected reality*, the (idealized) model of reality that is used by the PMS to reason, and the *physical reality*, the real world with the actual values of conditions and outcomes. The physical reality Φ_s reflects the concept of “now”, i.e., what is happening in the real environment whilst the process is under execution. In general, a task can only be performed in a given physical reality Φ_s if and only if that reality satisfies the *preconditions* Pre_i of that task. Moreover, each task has also a set of *effects* Eff_i that change the current physical reality Φ_s into a new physical reality Φ_{s+1} .

A PMS that takes in input such a process specification should guarantee that each task is executed correctly, i.e., with an output that satisfies the process specification itself. In fact, at execution time, the process can be easily invalidated because of task failures or since the environment may change due to some external event. For this purpose, the concept of *expected reality* Ψ_s is given. A recovery procedure is needed if the two realities are different from each other. An execution monitor is responsible for detecting whether the gap between the expected and physical realities is such that the original process δ_0 cannot progress its execution. In that case, the PMS has to find a recovery process δ_a that repairs δ_0 and removes the gap between the two kinds of reality.

We developed three different mechanisms for deriving a recovery procedure δ_a . Currently, the adaptation algorithm deployed in SmartPM synthesizes a linear process δ_a (i.e., a process consisting of a sequence of tasks) and inserts it at a given point of the original process - specifically, that point of the process where the deviation was first noted. To provide more details, let us assume that the current process is $\delta_0 = (\delta_1; \delta_2)$ in which δ_1 is the part of the process already executed and δ_2 is the

part of the process which remains to be executed when a deviation is identified. The adapted process is $\delta'_0 = (\delta_1; \delta_a; \delta_2)$. However, whenever a process needs to be adapted, every running task is interrupted, since the “repair” sequence of tasks $\delta_a = [t_1, \dots, t_n]$ is placed before them. Thus, active branches can only resume their execution after the repair sequence has been executed. This last requirement is fundamental to avoid the risk of introducing data inconsistencies during the repair phase.

3.1.1 Representing Tasks in SmartPM

An end user that wants to interact with the SmartPM prototype is not required to model a process and contextual data through complex languages such as Situation Calculus and IndiGolog. We use those languages for reasoning on the process under execution and for building automatically recovery procedures at run-time. In the following chapter, we describe how the end-user can build an admissible process specification for SmartPM with the SmartPM Definition Tool, that allows to graphically represent the control flow of a process through the BPMN language [143] and to annotate in a descriptive way (through the SmartML language) the tasks composing the process. Before to execute a process defined with BPMN and annotated with SmartML, the SmartPM system automatically translates those specifications in situation calculus and IndiGolog readable formats.

In SmartPM, tasks are collected in a specific repository, and each task can be considered as a single step that consumes input data and produces output data. For example, if we consider our case study introduced in Section 1.3, the task *go* can be described with SmartML as follows:

```
<task>
  <name>go</name>
  <parameters>
    <arg>from - Location_type</arg>
    <arg>to - Location_type</arg>
  </parameters>
  <precondition>at[SRVC] == from AND
    isConnected[SRVC]
  </precondition>
  <effects>
    <supposed>at[SRVC] = to</supposed>
  </effects>
</task>
```

The task *go* involves two input parameters *from* and *to* of type *Location*, representing the starting and arrival locations, while *SRVC* represents the service (i.e., the process participant) that will execute the task at run-time. An instance of this task can be executed only if the service *SRVC* is currently at the starting location *from* and is connected to the team’s network. The supposed consequence of task execution is that *SRVC* moves from the starting to the arrival location, and this is reflected by assigning to *at[SRVC]* the value *to* in the effect. But, in the real world, the execution of *go* can lead to a different outcome from the expected one. If *SRVC* reaches a location *to'* different from *to*, it means that the expected reality changes as the task was executed correctly, while the physical reality continues to reflect the

real world, by assuming that *SRVC* has reached *to'*. In those cases, SmartPM will be able to adapt the process for re-aligning the two realities without any pre-defined policy.

Let us consider that not every task has an outcome that requires the SmartPM system to monitor the two realities. This is the case of the task *move*:

```
<task>
  <name>move</name>
  <parameters>
    <arg>from - Location_type</arg>
    <arg>to - Location_type</arg>
  </parameters>
  <precondition>atRobot[SRVC] == from AND
                batteryLevel[SRVC]>=moveStep[] AND
                isRobotConnected[SRVC]
  </precondition>
  <effects>
    <supposed>atRobot[SRVC] = to</supposed>
    <automatic>batteryLevel[SRVC] -= moveStep[]</automatic>
  </effects>
</task>
```

A service *SRVC* that execute the task *move* (this task has been thought for representing the robots movement) should reach the location *to*, but the supposed effect involves a contextual property (i.e., *atRobot[SRVC]*, used for recording the current position of a robot) flagged at design-time as *not relevant*, meaning that we need only to update the physical reality with the real output returned by the task, but no update to the expected reality is required. The same is true for the *automatic* effect that decreases the battery level of the robot of a fixed quantity equal to *moveStep[]*. An automatic effect is automatically applied when the task completes without the need to consider the task outcomes.

A thorough description of the SmartML language is given in Section 4.3.1. In this chapter, we concentrate on modeling processes, data and tasks through situation calculus and IndiGolog, that allow us to perform some reasoning for automatically adapting faulty process instances at run-time.

3.1.2 Resource Model and Task Life-Cycle in SmartPM

When allocating tasks to the resources involved in business process execution, PMSs typically adopt a *pull-based* approach where the system offers each task to one or more resources qualified for it (e.g., using a *role-based* distribution approach) and a resource chooses among the offered items and commits to undertake the execution. This approach is motivated by the fact that, in business settings, a PMS has to leverage the interests, priorities, needs and constraints of single participants, and the overall revenue of the organizations they belong to.

The SmartPM prototype was instead initially developed for coordinating emergency management processes. Those dynamic processes are highly critical and time demanding, as well as they often need to be carried out within strictly specified deadlines. Therefore, it is inapplicable to use a *pull* mechanism for task assignment, as the risk is to have some task(s) waiting indefinitely for being chosen and executed.




Figure 3.1. The task-life cycle in SmartPM.

In emergency management, the personal benefit of single operators is unimportant with respect to the overall effectiveness of rescue operations. Consequently, it is preferable a *push-based* approach, where the system dynamically selects a resource qualified for executing a given task (e.g., using a *role-based* or a more specific *capability-based* distribution approach) and directly allocates the work item to the selected resource. According to this approach, in SmartPM each task is assigned **to only one resource**, and moreover, each resource gets assigned **at most one task**. In our model, resources involved in process executions are classified as either *human* (e.g., first responders that act on the field), or *non-human* (e.g., robots whose purpose is to maintain the network connection). However, we abstract all possible resources as *services*, i.e., entities capable of executing specific tasks [112].

The allocation of tasks to services and their execution during process enactment, as well as performed recovery and adaptation procedures, determine the life-cycle of a task, as illustrated in the state-transition diagram in Fig. 3.1. In SmartPM, each task making part of the process control flow that was never been assigned to any service lies in a *Not Assigned* state. A task becomes *Assigned* after the PMS assigns it to a specific service qualified for executing the task and upon fulfillment of control-flow preconditions. With respect to the general task life-cycle defined in [112], in our model a task is never *offered* to one or more services and is directly *assigned* to a service for execution. When a service is ready for executing the task, it notifies the SmartPM system through its *Task Handler* about the starting of the task itself; as a consequence the task moves to the *started* state. A started task may then be temporarily *suspended*, specifically when the system needs to restructure and modify the process for adaptation purposes, and successfully *completed*. A task failure due to the impossibility of the task complete its execution is seen by SmartPM as the task has completed with an “undefined” outcome.

3.2 Preliminaries

In this section we introduce some preliminary notions, namely *Situation Calculus* and *IndiGolog*, that are used as proper formalisms to reason about processes and exogenous events. Then, we will give some insights around the use of *classical planning techniques*. This section is not meant to give an all-comprehensive and very formal introduction of the notions. It aims mostly at giving an overall insight

to those who are not very expert on such topics.

3.2.1 Situation Calculus

The Situation Calculus [82, 107] is a logical language specifically designed for representing dynamically changing worlds in which all changes are the result of named *actions*. A possible history of actions is represented by a so-called *situation*, a first-order term in the language. The constant S_0 denotes the initial situation, where no actions have yet been performed. Sequences of actions are built using the special function symbol $do(a, s)$, that denotes the successor situation resulting from performing action a in situation s .

Example 3.2.1. Let us suppose we defined an action $open(x)$ concerning the opening of a door x . The situation term $do(open(d_2), do(open(d_1), S_0))$ denotes the situation resulting from first opening door d_1 in S_0 and then opening door d_2 .

In general, the Situation Calculus can be seen as a “dialect” of the first-order logic in which situations and actions are explicitly taken to be objects in the domain. Note that constants and function symbols for actions are completely application dependent. Properties that hold in a situation are called *fluents*. Technically, these are predicates taking a situation term as their last argument.

Example 3.2.2. The fluent $DoorOpen(x, s)$ may denote that door x is open in situation s . The formula $\neg DoorOpen(d_1, s) \wedge DoorOpen(d_1, do(open(d_1), s))$ says that the door d_1 is closed in some situation s , but it will be opened in the situation that results from opening it in s .

A distinguished predicate $Poss(a, s)$ is used to state that action a is executable in s .

Example 3.2.3. The predicate $Poss(open(x), S_0)$ says that is possible to open the door x in the initial situation S_0 .

To reason about a changing world, it is necessary to have beliefs not only about what is true initially, but also about how the world changes as result of actions. Actions typically have *preconditions*, that is, conditions that need to be true for the action to occur. Formulas like the one in the Example 3.2.4 are called *preconditions axioms*.

Example 3.2.4. The action of opening a door in situation s is possible only if the door is closed in s :

$$Poss(open(x), s) \equiv \neg DoorOpen(x, s).$$

Actions typically have also *effects*, that is, fluents that are changed as a result of performing the action.

Example 3.2.5. If a door x is closed and not locked, the action $open(x)$ can be used to open it:

$$(\neg DoorOpen(x, s) \wedge \neg Locked(x, s)) \supset DoorOpen(x, do(open(x), s)).$$

Example 3.2.6. To close a door causes it to be no more opened:

$$\neg DoorOpen(x, do(close(x), s)).$$

These formulas are called *effects axioms*; they can be *positive* (cf. Example 3.2.5) if they describe when a fluent becomes true, and *negative* (cf. Example 3.2.6) otherwise. In general, for any fluent $F(\vec{x}, s)$ we can rewrite all positive effects axioms as a single formula of the form:

$$P_F(\vec{x}, a, s) \supset F(\vec{x}, do(a, s)) \quad (3.1)$$

and all the negative effect axioms as a single formula of the form:

$$N_F(\vec{x}, a, s) \supset \neg F(\vec{x}, do(a, s)) \quad (3.2)$$

If we assume that formulas 3.1 and 3.2 characterize all the conditions under which an action a changes the value of a fluent F (this is called the *completeness assumption*, cf. [8]), and if we state that no action a satisfies the condition of making the fluent F both true and false at the same time, then all the effects axioms that are linked to a specific fluent F can be combined in a single *successor state axiom* that completely characterizes the value of F in the successor state resulting from performing action a in situation s .

$$F(\vec{x}, do(a, s)) \equiv P_F(\vec{x}, a, s) \vee (F(\vec{x}, s) \wedge \neg N_F(\vec{x}, a, s)) \quad (3.3)$$

Specifically, F is true after doing a if and only if before doing a , P_F (the positive effect condition for F) was true, or both F and $\neg N_F$ (the negative effect condition for F) were true.

Example 3.2.7. The successor state axiom for fluent $DoorOpen(x, s)$ is as follows:

$$\begin{aligned} DoorOpen(x, do(a, s)) \equiv \\ (a = open(x) \wedge \neg Locked(x, s) \wedge \neg DoorOpen(x, s)) \vee \\ DoorOpen(x, s) \wedge a \neq close(x); \end{aligned}$$

That is, a door is open after an action a has been performed in s iff a denotes the action of opening that door and the door is closed and not locked, or the door is already open in s and the action a is not that one of closing it.

The use of successor state axioms for capturing how a specific fluent changes after performing some action solves automatically the so-called *frame problem* [8]. In fact, to really know how the world can change, it should be also necessary to know what fluents are *unaffected* by performing an action.

Example 3.2.8. To open a door x does not change its color c :

$$\neg Color(x, c, s) \longrightarrow Color(x, c, do(open(x), s)).$$

These kinds of formulas are called *frame axioms*, because they limit or frame the effects of an action, and they would be required to fully capture the dynamism of a situation. Usually, it is expected that only a very small number of actions affects the value of a specific fluent; the rest leave it invariant (e.g., the color of a door

is unaffected by opening that same door). The *frame problem* forces to know and reason with a large number of frame axioms; if A is the number of actions and F the number of fluents, we should reason on about $2xAx F$ facts that do not change when we are performing an action. A simple solution to the frame problem has been proposed in [107, 8], and consists of using:

- precondition axioms, one per action (specifying when the action is executable);
- successor state axioms, one per fluent (capturing the effects and non-effects of actions);
- *initial state axioms* describing what is true initially (i.e., what is true in the initial situation S_0);
- unique name axioms for actions (stating that the only action terms that can be equal are two identical actions with identical arguments):
 - $A(\vec{x}) = A(\vec{y}) \supset (x_1 = y_1) \wedge \dots \wedge (x_n = y_n)$
 - $A(\vec{x}) \neq B(\vec{y})$ where A and B are distinct action names.

The completeness assumption for the effects of actions allows to conclude that action that are not mentioned explicitly in the effect axioms leave the fluent invariant (so, we do not need anymore to represent explicitly frame axioms).

Using the Situation Calculus

Given a knowledge base KB containing facts expressed in the situation calculus, there are various sorts of reasoning tasks we can consider [107, 8]. Two basic reasoning tasks are *projection* and *legality testing*.

Given a sequence of actions and some initial situation S_0 , the *projection task* aims at determining what would be true if those actions were performed starting in that initial situation. In [8], the problem is formalized as follows:

Projection Task Suppose that $\xi(s)$ is a formula with a single free variable s of the situation sort, and that \vec{a} is a sequence of actions $\langle a_1, \dots, a_n \rangle$. To find out if $\xi(s)$ would be true after performing \vec{a} starting in the initial situation S_0 , we determine whether or not $\text{KB} \models \xi(\text{do}(\vec{a}, S_0))$, where $\text{do}(\vec{a}, S_0)$ is an abbreviation for $\text{do}(a_n, \text{do}(a_{n-1}, \dots, \text{do}(a_2, \text{do}(a_1, S_0)) \dots))$, and for S_0 itself when $n = 0$.

Example 3.2.9. Let us suppose that in S_0 the doors d_1 and d_2 are closed and locked:

$$\text{Locked}(d_1, S_0) \wedge \text{Locked}(d_2, S_0) \wedge \neg \text{DoorOpen}(d_1, S_0) \wedge \neg \text{DoorOpen}(d_2, S_0)$$

and $\xi(s) \equiv (\text{DoorOpen}(d_1, s) \wedge \text{DoorOpen}(d_2, s))$. The sequence of actions:

$$\langle \text{unlock}(d_1), \text{open}(d_1), \text{unlock}(d_2), \text{open}(d_2) \rangle$$

makes the formula $\xi(s)$ equal to true. In other words, $\xi(s)$ holds in the situation:

$$s = \text{do}(\text{open}(d_2), \text{do}(\text{unlock}(d_2), \text{do}(\text{open}(d_1), \text{do}(\text{unlock}(d_1), S_0))))$$

Another interesting task performable through situation calculus is the *legality testing task*, that determines whether a sequence of actions leads to a legal situation. In [8], the problem is formalized as follows:

Legality Task Suppose that \vec{a} is a sequence of actions $\langle a_1, \dots, a_n \rangle$. To find out if \vec{a} can be legally performed starting in the initial situation S_0 , we determine whether or not $\text{KB} \models \text{Poss}(a_i, \text{do}(\langle a_1, \dots, a_{i-1} \rangle, S_0))$ for every i such that $1 \leq i \leq n$.

Example 3.2.10. For example, given the initial situation:

$$\text{Locked}(d_1, S_0) \wedge \text{Locked}(d_2, S_0) \wedge \neg \text{DoorOpen}(d_1, S_0) \wedge \neg \text{DoorOpen}(d_2, S_0)$$

although the situation term $\text{do}(\text{open}(d_2), \text{do}(\text{unlock}(d_1), S_0))$ is well-formed, it is not a legal situation, because the precondition for opening d_2 states that d_2 has to be unlocked, while the first action in the situation term unlocks door d_1 . The situation term:

$$s = \text{do}(\text{open}(d_2), \text{do}(\text{unlock}(d_2), \text{do}(\text{open}(d_1), \text{do}(\text{unlock}(d_1), S_0))))$$

is instead legal, since the sequence of actions can be legally performed starting from S_0 .

3.2.2 Indigolog

On top of Situation Calculus action theories, one can define complex control behaviors by means of high-level programs expressed in Golog-like programming languages. Specifically we focus on IndiGolog [25], which provides a set of programming constructs sufficient for defining every well-structured process as defined in [136].

IndiGolog is a programming language for autonomous agents that sense their environment and do planning as they operate. The programmer provides a high-level nondeterministic program involving domain-specific actions and tests to perform the agent's tasks. The IndiGolog interpreter then reasons about the preconditions and effects of the actions in the program to find a legal terminating execution. To support this, the programmer provides a declarative specification of the domain (i.e., primitive actions, preconditions and effects, what is known about the initial state) in the Situation Calculus. IndiGolog programs are executed **online** together with sensing the environment and monitoring for events.

IndiGolog derives from ConGolog to which it adds basically the lookahead *search operator*. Such operator allows to simulate the execution of a process with the aim of searching for a successful termination before actually performing the program in the real world. In its turn ConGolog [24] extends the original Golog [72] by introducing construct for current execution of different operations. Table 3.1 summarizes the constructs of IndiGolog used in this thesis.

In the first line, a stands for a situation calculus action term whereas, in the second line, $\sigma?$ stands for a formula over situation calculus predicates and fluents that needs to be evaluated when reached.

There are some constructs for concurrent programming. In particular $(\delta_1 \parallel \delta_2)$ expresses the concurrent execution (interpreted as interleaving) of the programs δ_1

Table 3.1. IndiGolog constructs.

Construct	Meaning
a	A primitive action.
$\sigma?$	Wait while the σ condition holds.
$(\delta_1; \delta_2)$	Sequence of two sub-programs δ_1 and δ_2 .
$proc P(\vec{v}) \delta$	Invocation of a IndiGolog procedure δ passing a vector \vec{v} of parameters.
$(\delta_1 \delta_2)$	Non-deterministic choice among (sub-)programs δ_1 and δ_2 .
$if \sigma then \delta_1 else \delta_2$	Conditional statement: if σ holds, δ_1 is executed; otherwise δ_2 .
$while \sigma do \delta$	Iterative invocation of δ .
$(\delta_1 \delta_2)$	Concurrent execution.
$(\delta_1 \gg \delta_2)$	Prioritized concurrent execution.
δ^*	Non-deterministic iteration of program execution.
$\pi a. \delta$	Non-deterministic choice of argument a followed by the execution of δ .
$\langle \sigma \rightarrow \delta \rangle$	δ is repeatedly executed until σ becomes false, releasing control to anyone else able to execute.
$\Sigma(\delta)$	search operator
$send(\Upsilon, \vec{v})$	a vector \vec{v} of parameters is passed to an external program Υ .
$receive(\Upsilon, \vec{v})$	a vector \vec{v} of parameters is received by an external program Υ .

and δ_2 . Another concurrent programming construct is $(\delta_1 \gg \delta_2)$, where δ_1 has higher priority than δ_2 , and δ_2 may only execute when δ_1 is done or blocked.

The list in Table 3.1 includes also some nondeterministic constructs. For example, $(\delta_1 | \delta_2)$ nondeterministically chooses between programs δ_1 and δ_2 . Test actions $\sigma?$ can be used to control which branches may be executed, e.g., $[(\sigma?; \delta_1) | (\neg\sigma?; \delta_2)]$ will perform δ_1 when σ is *true* and δ_2 when σ is *false* (we use $[...]$ and $(...)$ interchangeably to disambiguate structure in programs). The construct $\pi a. \delta$, nondeterministically picks a binding for the variable a and performs the program δ for this binding of a . Finally, δ^* , performs δ zero or more times. $\pi a_1, \dots, a_n. \delta$ is an abbreviation for $\pi a_1. \dots. \pi a_n. \delta$.

The constructs **if σ then δ_1 else δ_2** and **while σ do δ** are the synchronized versions of the usual if-then-else and while-loop. They are synchronized in the sense that testing the condition σ does not involve a transition per se: the evaluation of the condition and the first action of the branch chosen are executed as an atomic unit. So these constructs behave in a similar way to the test-and-set atomic instructions used to build semaphores in concurrent programming.

Let's focus on the interrupt construct:

$$\langle \sigma \rightarrow \delta \rangle \stackrel{\text{def}}{=} \mathbf{while} \text{Interrupts_running} \mathbf{do}$$

$$\mathbf{if} \sigma \mathbf{then} \delta \mathbf{else} \text{false} \mathbf{endIf}$$

$$\mathbf{endWhile}$$

To see how this works, first assume that the special fluent *Interrupts_running* is identically true. When an interrupt $\langle \sigma \rightarrow \delta \rangle$ gets control from higher priority processes, suspending any lower priority processes that may have been advancing, it repeatedly executes δ until σ becomes false. Once the interrupt body δ completes its execution, the suspended lower priority processes may resume. The control release also occurs if σ cannot progress (e.g., since no action meets its precondition).

Since IndiGolog provides flexible mechanisms for interfacing with other programming languages such as Java or C and for socket communication (cf. Section 4.2), for

our convenience we have defined here two more abstract constructs to send/receive parameters as well as values with external programs, defined out of the range of the IndiGolog process.

From a formal point of view, two predicates are introduced to specify program transitions:

- $Trans(\delta', s', \delta'', s'')$, given a program δ' and a situation s' , returns (i) a new situation s'' resulting from executing a single step of δ' , and (ii) δ'' which is the remaining program to be executed.
- $Final(\delta', s')$ returns true when the program δ' can be considered successfully completed in situation s' , i.e., δ' is *legally terminated* in situation s' .

The predicates $Trans$ and $Finals$ for programs are characterized by the following set of axioms:

1. Empty program:

$$Trans(nil, s, \delta', s') \Leftrightarrow \text{false.}$$

$$Final(nil, s) \equiv \text{True.}$$

2. Primitive actions:

$$Trans(a, s, \delta', s') \Leftrightarrow Poss(a[s], s) \wedge \delta' = nil \wedge s' = do(a[s], s).$$

$$Final(a, s) \equiv \text{false.}$$

3. Test/wait actions:

$$Trans(\sigma?, s, \delta', s') \Leftrightarrow \sigma[s] \wedge \delta' = nil \wedge s' = s.$$

$$Final(\sigma?, s) \equiv \text{false.}$$

4. Sequence:

$$Trans(\delta_1; \delta_2, s, \delta', s') \Leftrightarrow \\ \exists \gamma. \delta' = (\gamma; \delta_2) \wedge Trans(\delta_1, s, \gamma, s') \vee Final(\delta_1, s) \wedge Trans(\delta_2, s, \delta', s').$$

$$Final(\delta_1; \delta_2, s) \equiv Final(\delta_1, s) \wedge Final(\delta_2, s).$$

5. Nondeterministic branch:

$$Trans(\delta_1 \mid \delta_2, s, \delta', s') \Leftrightarrow Trans(\delta_1, s, \delta', s') \vee Trans(\delta_2, s, \delta', s').$$

$$Final(\delta_1 \mid \delta_2, s) \equiv Final(\delta_1, s) \vee Final(\delta_2, s).$$

6. Nondeterministic choice of argument:

$$Trans(\pi v. \delta, s, \delta', s') \Leftrightarrow \exists x. Trans(\delta_x^v, s, \delta', s').$$

$$Final(\pi v. \delta, s) \equiv \exists x. Final(\delta_x^v, s).$$

7. Nondeterministic iteration:

$$\begin{aligned} \text{Trans}(\delta^*, s, \delta', s') &\Leftrightarrow \exists \gamma. (\delta' = \gamma; \delta^*) \wedge \text{Trans}(\delta, s, \gamma, s'). \\ \text{Final}(\delta^*, s) &\equiv \text{True}. \end{aligned}$$

8. Synchronized conditional:

$$\begin{aligned} \text{Trans}(\mathbf{if} \sigma \mathbf{then} \delta_1 \mathbf{else} \delta_2 \mathbf{endIf}, s, \delta', s') &\Leftrightarrow \\ &\sigma[s] \wedge \text{Trans}(\delta_1, s, \delta', s') \vee \neg\sigma[s] \wedge \text{Trans}(\delta_2, s, \delta', s'). \\ \text{Final}(\mathbf{if} \sigma \mathbf{then} \delta_1 \mathbf{else} \delta_2 \mathbf{endIf}, s) &\equiv \\ &\sigma[s] \wedge \text{Final}(\delta_1, s) \vee \neg\sigma[s] \wedge \text{Final}(\delta_2, s). \end{aligned}$$

9. Synchronized loop:

$$\begin{aligned} \text{Trans}(\mathbf{while} \sigma \mathbf{do} \delta \mathbf{endWhile}, s, \delta', s') &\Leftrightarrow \\ &\exists \gamma. (\delta' = \gamma; \mathbf{while} \sigma \mathbf{do} \delta) \wedge \sigma[s] \wedge \text{Trans}(\delta, s, \gamma, s'). \\ \text{Final}(\mathbf{while} \sigma \mathbf{do} \delta \mathbf{endWhile}, s) &\equiv \neg\sigma[s] \vee \text{Final}(\delta, s). \end{aligned}$$

10. Concurrent execution:

$$\begin{aligned} \text{Trans}(\delta_1 \parallel \delta_2, s, \delta', s') &\Leftrightarrow \\ &\exists \gamma. \delta' = (\gamma \parallel \delta_2) \wedge \text{Trans}(\delta_1, s, \gamma, s') \vee \exists \gamma. \delta' = (\delta_1 \parallel \gamma) \wedge \text{Trans}(\delta_2, s, \gamma, s'). \\ \text{Final}(\delta_1 \parallel \delta_2, s) &\equiv \text{Final}(\delta_1, s) \wedge \text{Final}(\delta_2, s). \end{aligned}$$

11. Prioritized concurrency:

$$\begin{aligned} \text{Trans}(\delta_1 \gg \delta_2, s, \delta', s') &\Leftrightarrow \\ &\exists \gamma. \delta' = (\gamma \gg \delta_2) \wedge \text{Trans}(\delta_1, s, \gamma, s') \vee \\ &\exists \gamma. \delta' = (\delta_1 \gg \gamma) \wedge \text{Trans}(\delta_2, s, \gamma, s') \wedge \neg \exists \zeta, s''. \text{Trans}(\delta_1, s, \zeta, s''). \\ \text{Final}(\delta_1 \gg \delta_2, s) &\equiv \text{Final}(\delta_1, s) \wedge \text{Final}(\delta_2, s). \end{aligned}$$

12. Concurrent iteration:

$$\begin{aligned} \text{Trans}(\delta^\parallel, s, \delta', s') &\Leftrightarrow \exists \gamma. \delta' = (\gamma \parallel \delta^\parallel) \wedge \text{Trans}(\delta, s, \gamma, s'). \\ \text{Final}(\delta^\parallel, s) &\equiv \text{True}. \end{aligned}$$

The **off-line** execution of programs, which is the kind of execution originally proposed for Golog [72] and ConGolog [24] is characterized using the $Do(\delta, s, s')$ predicate, which means that there is an execution of program δ that starts in situation s and terminates in situation s' :

$$Do(\delta, s, s') \Leftrightarrow \exists \delta'. \text{Trans}^*(\delta, s, \delta', s') \wedge \text{Final}(\delta', s')$$

where Trans^* is the definition of the reflective and transitive closure of Trans . Thus there is an execution of program δ that starts in situation s and terminates in

situation s' if and only if we can perform 0 or more transitions from program δ in situation s to reach situation s' with program δ' remaining, at which point one may legally terminate. Notice that there may be more than one resulting situation s' since IndiGolog programs can be non-deterministic (e.g., due to concurrency).

The off-line execution model of Golog and ConGolog requires the executor to search over the whole program to find a complete execution before performing any action. This is obviously problematic for agents that need to sense their environment as they operate. On the contrary, the strength of IndiGolog is that it provides an **online** execution model that allows to execute actions on the real world, to update its knowledge after each action execution and to monitor for possible exogenous events or actions not executed as expected.

Finally, to cope with the impossibility of backtracking actions executed in the real world, IndiGolog incorporates a new programming construct, namely the *search operator* $\Sigma(\delta)$, which is used to specify that lookahead should be performed over the (nondeterministic) program δ to ensure that nondeterministic choices are resolved in a way that guarantees its successful completion. More precisely, let δ be any IndiGolog program, which provides different alternative executable actions. When the interpreter encounters program $\Sigma(\delta)$, before choosing among alternative executable actions of δ , it performs reasoning in order to decide for a step which still allows the rest of δ to terminate successfully. Formally, according to [26], the semantics of the search operator is:

$$Trans(\Sigma(\delta), s, \Sigma(\delta'), s') \Leftrightarrow Trans(\delta, s, \delta', s') \wedge \exists s^*. Do(\delta', s', s^*).$$

If δ is the entire program under consideration, $\Sigma(\delta)$ emulates completely its off-line execution as if it is executed on-line.

3.2.3 Classical Planning

Automated Planning [90] is the branch of AI that consists in the deliberation process of building plans, i.e., organized actions, in order to fulfill some pre-stated objectives. Typically, these plans are executed by intelligent agents, and the solution amounts to synthesize agent's plans satisfying a goal specification.

There exist several forms of planning. We can distinguish between *classical planning* domains, *conditional planning* and *conformant planning*. *Classical planning* domains are fully observable, static, and deterministic, in which plans can be computed in advance and then applied unconditionally. Therefore a *plan* is a simple sequence $a_0, a_1, \dots, a_n \in Act^*$ of actions, for a fixed action alphabet Act . *Conditional* (or *contingency*) *planning* deals instead with bounded indeterminacy by constructing a conditional plan with different branches for the different contingencies that may arise, and even though plans are pre-computed, the agent finds out which part of the plan to execute by including sensing actions in the plan to test for the appropriate conditions. Therefore, conditional plans can be thought of as tree-like structures, in contrast with sequential plans that are instead action sequences. Finally, *conformant planning* [120] aims to construct standard, sequential plans that are executed in partially-observable settings, without perception. Namely, they are required to achieve the goal in all possible circumstances, regardless of the true initial state and the actual action outcomes.

In this thesis, we focus on classical planning techniques. Classical planning has made huge advances in the last twenty years, leading to solvers able to create plans with thousands of actions for problems described by hundreds of propositions. The standard representation language for classical planners is known as the Planning Domain Definition Language (PDDL (cf. [46])); it allows one to formulate a problem PR through the description of the initial state of the world $init_{PR}$, the description of the desired goal condition $goal_{PR}$ and a set of possible actions. An action definition defines the conditions under which an action can be executed, called *preconditions*, and its *effects* on the state of the world. The set of all action definitions Ω is the *domain* PD of the planning problem. Each action $a \in \Omega$ has a precondition list and an effect list, denoted respectively as Pre_a and Eff_a . For example, let us consider the following (trivial) planning domain named “example”:

```
(define (domain example)
  (:predicates
    (x) (y) (z)
    (k) (v) (s)
  )
  (:action t1
    :precondition (x)
    :effect (and (not(x) (z) (k))))
  (:action t2
    :precondition (z)
    :effect (y))
  (:action t3
    :precondition (v)
    :effect (s))
  (:action t4
    :precondition (k)
    :effect (x))
)
```

The meaning is straightforward. There are 6 predicates and 4 planning actions, and for each action some preconditions and effects are defined. For example, the first action definition states that for executing t_1 , the predicate x must hold. Then, it states that a successful execution of t_1 guarantees that predicates $\neg x$, z and k will hold together. Given a specific planning domain, it is possible to define a proper planning problem:

```
(define (problem pr) (:domain example)
  (:init
    (x) (v)
  )
  (:goal
    (and
      (y) (s)
    )
  )
)
```

It states that in $init_{PR}$ only predicates x and v are known to be *true*, and the goal condition $goal_{PR}$ is a formula where the conjunction of y and s is *true*. Note that one of the main assumption we rely in building planning problem is the well known *closed-world assumption* [106]. Basically, it states that all unmentioned literals in a

planning problem are considered to be *false*. Hence, for example, in the planning problem defined above we are supposing that predicates z , k , y and s are initially *false*.

A planner that works on such inputs generates a sequence of actions (the *plan*) that corresponds to a path from the initial state to a state meeting the goal condition. A simple plan satisfying the above planning problem consists first in executing action $t3$ (since predicate v is *true* in the initial state), that makes s equal to *true*, and then on executing $t1$ and $t2$ in sequence, for turning the value of y to *true*. This simple sequence composed by three actions is an *optimal solution*, since it requires the minimum number of steps for satisfying the goal condition. Other possible plans that satisfy the goal condition but require more steps if compared with the optimal solution, are called *suboptimal*.

In this chapter, we represent planning domains and planning problems making use of PDDL version 2.2 [37]. PDDL 2.2 is characterized for enabling the representation of realistic planning domains, which include (in particular) actions and goals involving numerical expressions, operators with universally quantified effects or existentially quantified preconditions, operators with disjunctive or implicative preconditions, derived predicates and plan metrics. However, currently, our formalism does not allow to represent conditional and universally quantified effects.

Specifically, we synthesize our recovery plan through the LPG-td planner [45] (Local search for Planning Graphs). It is a planner based on local search and planning graphs that handles PDDL 2.2 domains. The basic search scheme of LPG-td is inspired by Walksat [118], an efficient procedure to solve SAT-problems. The search space of LPG consists of “action graphs”, particular subgraphs of the planning graph representing partial plans. The search steps are certain graph modifications transforming an action graph into another one. LPG-td exploits a compact representation of the planning graph to define the search neighborhood and to evaluate its elements using a parametrized function, where the parameters weight different types of inconsistencies in the current partial plan, and are dynamically evaluated during search using discrete Lagrange multipliers. The evaluation function uses some heuristics to estimate the “search cost” and the “execution cost” of achieving a (possibly numeric) precondition. The planner can produce good quality plans in terms of one or more criteria. This is achieved by an anytime process producing a sequence of plans, each of which is an improvement of the previous ones in terms of its quality. More details on the search algorithm and heuristics devised for this planner can be found at [45, 44].

3.3 Formalizing processes in IndiGolog

Our approach to process management relies on the IndiGolog platform [25], that is able to reason about a dynamic changing world by updating its knowledge after each action execution. The PMS provided by the SmartPM system has been developed on top of the IndiGolog platform (in the rest of this thesis, when we refer to the IndiGolog PMS, or to the IndiGolog engine, we always mean the PMS developed through the IndiGolog platform). The IndiGolog PMS takes in input a theory of actions - specified in the situation calculus - representing the contextual environment in which the

process operates, and an IndiGolog program (specified through the constructs shown in Table 3.1) that reflects the control flow of the dynamic process to be executed. To this end, in this section, we describe how processes can be formalized in situation calculus and IndiGolog. We refer to our case study introduced in Section 1.3.

To denote the various objects of interest, we make use of the following domain-independent predicates (that is, non-fluent rigid predicates):

- $Service(c)$: c is a service (i.e., a process participant);
- $Task(t)$: t is a task;
- $Capability(b)$: b is a capability;
- $Provides(c, b)$: service c provides the capability b ;
- $Requires(t, b)$: task t requires the capability b .

Example 3.3.1. Let us consider the case study defined in Section 1.3. We need to define 6 services (4 human actors and 2 robots), a repository of emergency-management tasks and a number of capabilities to be associated both to services and tasks.

$Service(act1)$, $Service(act2)$, $Service(act3)$, $Service(act4)$, $Service(rb1)$, $Service(rb2)$.

$Capability(hatchet)$, $Capability(extinguisher)$, $Capability(camera)$, $Capability(battery)$, $Capability(movement)$, $Capability(gprs)$, $Capability(digger)$, $Capability(powerpack)$.

$Task(chargebattery)$, $Task(go)$, $Task(move)$, $Task(extinguishfire)$, $Task(evacuate)$, $Task(removedebris)$, $Task(takephoto)$, $Task(updatestatus)$.

$Provides(act1, movement)$, $Provides(act1, gprs)$, $Provides(act1, extinguisher)$, $Provides(act1, camera)$, $Provides(act2, movement)$, $Provides(act2, gprs)$, $Provides(act2, hatchet)$, $Provides(act3, movement)$, $Provides(act3, gprs)$, $Provides(act3, hatchet)$, $Provides(act4, movement)$, $Provides(act4, powerpack)$, $Provides(act4, gprs)$, $Provides(rb1, battery)$, $Provides(rb1, digger)$, $Provides(rb2, battery)$, $Provides(rb2, digger)$.

$Requires(go, movement)$, $Requires(evacuate, hatchet)$, $Requires(takephoto, camera)$, $Requires(updatestatus, gprs)$, $Requires(extinguishfire, extinguisher)$, $Requires(move, battery)$, $Requires(removedebris, digger)$, $Requires(chargebattery, powerpack)$.

To refer to the ability of a service c to perform a certain task t , we introduce the following abbreviation:

$$Capable(c, t) \stackrel{\text{def}}{=} \forall b. Requires(t, b) \Rightarrow Provides(c, b). \quad (3.4)$$

That is, service c can carry out a certain task t if and only if c provides all capabilities required by the task t .

The life-cycle of a task involves the execution of four basic actions:

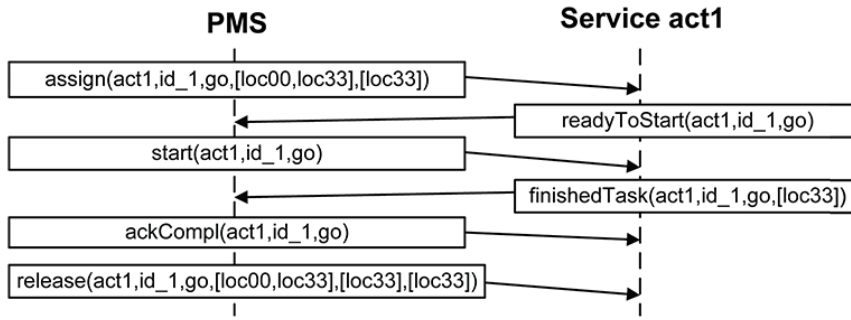


Figure 3.2. The protocol for task assignment and execution.

- $assign(c, id, t, \vec{i}, \vec{p})$: a task t with input $\vec{i} = [i_1, \dots, i_z]$ is assigned to a service c . With $\vec{p} = [p_1, \dots, p_w]$ we denote the list of *expected outputs* that t is supposed to return if its execution is successful;
- $start(c, id, t)$: service c is notified to start task t ;
- $ackCompl(c, id, t)$: service c acknowledges of the completion of task t ;
- $release(c, id, t, \vec{i}, \vec{p}, \vec{q})$: service c releases task t , that was executed with the list of inputs \vec{i} and expected outputs \vec{p} , and returns a list of *physical outputs* $\vec{q} = [q_1, \dots, q_w]$.

Every running task will be associated with a unique identifier id (to this end, we provide a list of admissible identifiers and a further predicate $Identifier(id)$ that is *true* if id represents a valid identifier) and with a number of input and output parameters. In this sense, the terms \vec{i} , \vec{p} and \vec{q} denote arbitrary sets of input/output, which depend on the specific task. The special constant \emptyset denotes empty input or output. Note that we suppose to work with domains in which services, tasks, input and output parameters are finite.

The actions performed by the process need to be “complemented” by other actions executed by the services themselves. The following are used to inform the PMS engine about how tasks execution is progressing:

- $readyToStart(c, id, t)$: service c declares to be ready to start performing t ;
- $finishedTask(c, id, t, \vec{q})$: service c declares to have completed the execution of task t with a list of physical outputs denoted by \vec{q} .

The protocol for a successful execution of a specific task t_j can be described as follows. After the PMS assigns t_j to a service c through the action $assign(c, id, t_j, \vec{i}, \vec{p})$, the service c has to report to be ready to execute the task itself by launching the action $readyToStart(c, id, t_j)$. When it does happen, the PMS can eventually perform the action $start(c, id, t_j)$, meaning that c has been formally allowed to start executing t_j . When c completes the execution of t_j , it invokes the action $finishedTask(c, id, t_j, \vec{q})$, with \vec{q} representing the real outcomes (we also call them “physical” outcomes) of t_j applied on the contextual environment. At this point the PMS can update the properties (i.e., the situation calculus fluents, see later) reflecting the evolution of the contextual scenario in which the process is under execution. In addition, the PMS first acknowledges the completion of t_j through the action $ackCompl(c, id, t_j)$, and then

releases the service c from the task t_j . After the execution of $release(c, id, t_j, \vec{i}, \vec{p}, \vec{q})$, service c is again ready for a new task assignment. In Fig. 3.2 we describe the protocol for assigning the task go to service $act1$ and executing it properly. In fact, when $act1$ launches the $finishedTask$ action (meaning the task has been completed by $act1$) it returns an expected output equal to the desired destination expected at design time. On the contrary, note that the scenario described in Fig. 1.3 reflects a wrong execution of go , with an output different from the one expected.

In order to describe the contextual scenario in which the process will be executed and to bound task inputs and outputs, there is the need to define some basic data types. The IndiGolog language already provides the *Boolean_type* (for representing boolean values) and the *Integer_type* (for representing integers from 0 to a specific customizable maximum value - the default is 30). Furthermore, to clearly delineating our case study, we also need to define a data type *Location_type* that represents locations in the area and a data type *Status_type* for describing if a specific location is on fire or buried by debris.

Example 3.3.2. The list of required data types for our case study is as follows:

```

Boolean_type(true).
Boolean_type(false).
Integer_type(0).
...
Integer_type(30).
Status_type(ok).
Status_type(fire).
Status_type(debris).
Location_type(loc00).
...
Location_type(loc33).

```

Data types can be associated to tasks for bounding input and output parameters. For this aim, we introduce the concept of *workitem*, that reflects the run-time container of a task, including the task identifier and its input/output parameters. This means that the same task can be assigned multiple times to services if and only if it is associated with different identifiers, i.e, if it is contained in different workitems. Before to execute the process, we need to clearly state which workitems are admissible for being executed by the PMS.

Example 3.3.3. The task go can be executed by the PMS if and only if it is associated to a couple of inputs $from$ and to , that represent starting and arrival locations, and to an expected output to , that corresponds to the expected destination to be reached after the task execution. Note that in this specific case the expected output corresponds to the second argument given as input to the task go .

$$\forall (from, to, id) \text{ s.t. } Identifier(id) \wedge Location_type(from) \wedge Location_type(to) \Rightarrow workitem(go, id, [from, to], [to]).$$

The formalization of processes in IndiGolog requires two distinct sets of fluents. A first set includes those *domain-independent fluents* that the PMS uses to manage the task life-cycle and the resource perspective of a process. Domain-independent fluents do not change across different domains. The second set concerns those fluents that we use to denote the data needed by a process instance and the properties of the contextual scenario in which the process is under execution. We call them *data fluents*, and their definition depends strictly on the specific process domain of interest.

Representing domain-independent fluents.

Task assignment is driven by the fluent $Free(c, s)$, which denotes whether a service c is available for task assignments in situation s . Basically, $Free(c, s)$ holds if the service c is free from any assignment in situation s . The corresponding successor state axiom can be defined as follows:

$$\begin{aligned} Free(c, do(a, s)) \Leftrightarrow & \\ & (\exists t, id, \vec{i}, \vec{p}, \vec{q} \text{ s.t. } a = \text{release}(c, id, t, \vec{i}, \vec{p}, \vec{q})) \vee \\ & (Free(c, s) \wedge \nexists t, id, \vec{i}, \vec{p} \text{ s.t. } a = \text{assign}(c, id, t, \vec{i}, \vec{p})) \end{aligned} \quad (3.5)$$

Therefore, service c is considered as free in the current situation if and only if it has just been released from a task assignment or it was free in the previous situation and no task has been assigned to it.

The domain-independent fluent $Assigned(c, id, t, s)$ aims at representing if a task t with identifier id has been assigned to a service c in situation s .

$$\begin{aligned} Assigned(c, id, t, do(a, s)) \Leftrightarrow & \\ & (\exists \vec{i}, \vec{p} \text{ s.t. } a = \text{assign}(c, id, t, \vec{i}, \vec{p})) \vee \\ & (Assigned(c, id, t, s) \wedge \nexists \vec{i}, \vec{p}, \vec{q} \text{ s.t. } a = \text{release}(c, id, t, \vec{i}, \vec{p}, \vec{q})) \end{aligned} \quad (3.6)$$

The fluent holds when a task t with identifier id is assigned to a service c , or when t has already been assigned to c , and no release action concerning t and c is launched in the current situation s .

A third fluent used for constraining the protocol for task assignment and execution is $Reserved$. The corresponding successor-state axiom can be defined as follows:

$$\begin{aligned} Reserved(c, id, t, do(a, s)) \Leftrightarrow & \\ & (Assigned(c, id, t, s) \wedge a = \text{readyToStart}(c, id, t)) \vee \\ & (Reserved(c, id, t, s) \wedge \nexists \vec{q} \text{ s.t. } a = \text{finishedTask}(c, id, t, \vec{q})) \end{aligned} \quad (3.7)$$

This fluent holds when a service c is ready for start executing a task t with identifier id that was already assigned to it, or if c is still executing t and no *finishedTask* action has been launched from c .

The *Reserved* fluent can be complemented with the *precondition axioms* defined for the *start* and *ackCompl* actions:

$$\begin{aligned} Poss(\text{start}(c, id, t), s) &\Leftrightarrow Reserved(c, id, t, s) \\ Poss(\text{ackCompl}(c, id, t), s) &\Leftrightarrow \neg Reserved(c, id, t, s) \end{aligned} \quad (3.8)$$

The meaning of those axioms is straightforward: $\text{start}(c, id, t)$ can be launched by the PMS only if $Reserved(c, id, t, s)$ is true, therefore after c has performed the

action $readyToStart(c, id, t)$ (cf. Equation 3.7). The action $ackCompl(c, id, t)$ can be instead invoked when $Reserved(c, id, t, s)$ does not hold anymore, so when c launches $finishedTask(c, id, t, \vec{q})$ (cf. again Equation 3.7). The combination of precondition axioms defined in Equation 3.8 and of the state-successor axiom represented in Equation 3.7 states that $ackCompl(c, id, t)$ never comes before $start(c, id, t)$, by guaranteeing the synchronization of the single steps provided by the protocol for task assignment and execution.

Representing data fluents.

Data Fluents need to be customized for every domain, and they can be used to constrain the task assignment and as guards into the expressions at decision points (e.g., for cycles, conditional statements). However, a data fluent X_φ ¹ is mainly meant to capture one of the outcomes of a (specific) task T . The general way for expressing a successor-state axiom of a data fluent is as follows:

$$\begin{aligned}
X_\varphi(\vec{\sigma}, do(a, s)) = q_j \equiv & \\
& (\exists c, id, \vec{i}, \vec{p}, \vec{q} \text{ s.t. } a = \text{release}(c, id, T, \vec{i}, \vec{p}, \vec{q})) \wedge \vec{q} = [q_0, \dots, q_j, \dots, q_w] \vee \\
& (X_\varphi(\vec{\sigma}, s) = q_j \wedge \\
& \neg \exists c, id, \vec{i}, \vec{p}, \vec{q}' \text{ s.t. } a = \text{release}(c, id, T, \vec{i}, \vec{p}, \vec{q}') \wedge \vec{q}' = [q'_0, \dots, q'_j, \dots, q'_w] \wedge (q'_j \neq q_j)).
\end{aligned} \tag{3.9}$$

The value of X_φ is changed to value q_j when the task T finishes with output $\vec{q} = [q_0, \dots, q_j, \dots, q_w]$ or, more formally, when the PMS launches the action $\text{release}(c, id, T, \vec{i}, \vec{p}, \vec{q})$, with $q_j \in \vec{q}$. With $\vec{\sigma}$ we denote the list (possibly empty) of arguments of a data fluent, which can be customized in according to the process needs. For example, one can require some fluents defined for each service c in the formalization:

Example 3.3.4. The task *go* is used for instructing a service (specifically, a human actor) *act* to move from a location *from* to a location *to*. When *act* terminates the execution of the task *go* that was assigned to her/him with identifier id_1 , its final position q_j in situation s is stored in a specific data fluent At_φ .

$$\begin{aligned}
At_\varphi(act, do(a, s)) = q_j \equiv & \\
& (\exists q_j \text{ s.t. } Location_type(q_j) \wedge \\
& a = \text{release}(act, id_1, go, [from, to], [to], [q_j])) \vee \\
& (At_\varphi(act, s) = q_j \wedge \\
& \neg \exists q'_j \text{ s.t. } Location_type(q'_j) \wedge \\
& a = \text{release}(act, id_1, go, [from, to], [to], [q'_j]) \wedge (q'_j \neq q_j)).
\end{aligned} \tag{3.10}$$

Example 3.3.5. The task *move* is used for instructing a robot *rb* to move from a location *from* to a location *to*. When *rb* terminates the execution of the task *move* that was assigned to it with identifier id_2 , its final position q_j in situation s is stored in a specific data fluent $AtRobot_\varphi$.

¹Sometimes we use *arguments-suppressed* formulas; these are uniform formulas with all arguments suppressed (e.g. X_φ denotes the arguments-suppressed expression for $X_\varphi(\vec{\sigma}, s)$).

$$\begin{aligned}
AtRobot_{\varphi}(rb, do(a, s)) = q_j &\equiv \\
(\exists q_j \text{ s.t. } Location_type(q_j) \wedge & \\
a = release(rb, id_2, move, [from, to], [to], [q_j])) \vee & \\
(AtRobot_{\varphi}(rb, s) = q_j \wedge & \\
\neg \exists q'_j \text{ s.t. } Location_type(q'_j) \wedge & \\
a = release(rb, id_2, move, [from, to], [to], [q'_j]) \wedge (q'_j \neq q_j)). &
\end{aligned} \tag{3.11}$$

Data fluents are mainly used for recording the effects of a task. If we analyze the descriptions of the tasks *go* and *move* provided in Section 3.1.1, we can see that the content of the $\langle effects \rangle$ tag coincides with the successor state axioms described above.

When a task returns some real-world outcome after its completion, we define that outcome as *supposed*, since its physical value may be different from the expected one as thought at design-time. However, sometimes it may happen that a task effect is *automatic*, i.e., it is applied every time a task completes its execution, independently by the outcomes returned by the task itself.

Example 3.3.6. After a robot *rb* completes the execution of the task *move* with identifier id_2 , its battery level decreases of a fixed quantity equal to $MoveStep_{\varphi}(s)$. The battery charge level of a robot can be represented with a fluent $BatteryLevel_{\varphi}$:

$$\begin{aligned}
BatteryLevel_{\varphi}(rb, do(a, s)) = v &\equiv \\
(\exists \vec{i}, \vec{p}, \vec{q} \text{ s.t. } a = release(rb, id_2, move, \vec{i}, \vec{p}, \vec{q}) \wedge & \\
v = BatteryLevel_{\varphi}(rb, s) - MoveStep_{\varphi}(s) \wedge Integer_type(v)) \vee & \\
(BatteryLevel_{\varphi}(rb, s) = v \wedge & \\
\neg \exists \vec{i}, \vec{p}, \vec{q} \text{ s.t. } a = release(rb, id_2, move, \vec{i}, \vec{p}, \vec{q})). &
\end{aligned} \tag{3.12}$$

The fluent holds every time a *release* action involving the task *move* and the service *rb* is launched by the PMS, and it results in decreasing the level of the battery of the fixed quantity $MoveStep_{\varphi}(s)$ by obtaining a new battery charge level equal to v .

The representation of a dynamic scenario may require to define some *fixed* data fluents that reflect the contextual properties of the scenario itself. For example, our case study requires to define a data fluent $Neigh(loc1, loc2, s)$ that expresses the neighborhood property between two locations *loc1* and *loc2*, and a fluent $Covered(loc, s)$ that indicates which locations *loc* are directly covered by the radio range provided by the main antenna. Moreover, some data fluents can be used for recording constant values. For example, the fluent $MoveStep(s)$ indicates the fixed quantity of battery that will be consumed by a robot after the execution of a *move* task. Those kinds of fluents, in general, never change their values during process execution.

Representing complex formulae and exogenous events.

The contextual scenario presented in the case study in Section 1.3 requires also to define some properties that help to understand if a service is connected to the network provided by the main antenna. To this end, we can express some situation calculus *abbreviations* for verifying the connection status of each service in every possible situation s . An abbreviation is a complex formula that does not depend directly from tasks effects, and that can be evaluated in each situation s . Specifically,

given a human service act , the abbreviation $isConnected(act, s)$ holds in situation s if the service act is situated in a location covered by the main network. Otherwise, if act is outside the main network, it results as being connected only if it is close (cf. fluent $Neigh$) to a location where a robot rb is situated, and the robot is - in its turn - connected to the network.

$$\begin{aligned}
isConnected(act, s) \equiv & \\
& provides(act, movement) \wedge At_{\varphi}(act, s) = loc \wedge Location_type(loc) \wedge \\
& (Covered(loc, s) \vee \\
& (\exists rb \text{ s.t. } provides(rb, battery) \wedge AtRobot(rb, s) = loc_{rb} \wedge \\
& Location_type(loc_{rb}) \wedge Neigh(loc, loc_{rb}, s) \wedge isRobotConnected(rb, s))).
\end{aligned} \tag{3.13}$$

We omit here the definition of the second abbreviation $isRobotConnected$ (interesting readers can find it in the appendix), that holds if a robot is connected to the main network on the basis of the connection policy described in Section 1.3.

Abbreviations can be used together with the data fluents for defining the *preconditions* of each task. In *IndiGolog*, we make explicit task preconditions through the axiom $Poss$ applied on the $assign$ basic action:

$$\begin{aligned}
Poss(assign(c, id, t, \vec{i}, \vec{p}), s) \Leftrightarrow & workitem(t, id, \vec{i}, \vec{p}) \wedge \\
& (X_{\varphi,1} \wedge \dots \wedge X_{\varphi,m}) \wedge (...list\ of\ abbreviations...).
\end{aligned}$$

A task t with input \vec{i} and expected outputs \vec{p} can be assigned to a service c iff the task t combined with the associated inputs/outputs represents an admissible workitem. Moreover, values of data fluents $X_{\varphi,y}$ (where y ranges over $\{1..m\}$) and abbreviations possibly included in the axiom should be evaluated.

Example 3.3.7. A task go with inputs $from$ and to can be assigned to a service act in situation s if and only if act is located in $from$ and is connected to the network.

$$\begin{aligned}
Poss(assign(act, id, go, [from, to], [to]), s) \Leftrightarrow & \\
& workitem(go, id, [from, to], [to]) \wedge At_{\varphi}(act, s) = from \wedge isConnected(act, s).
\end{aligned} \tag{3.14}$$

Finally, we conclude the section by pointing out that in dynamic domains it is typical that variables asynchronously change their value in an impredicative fashion. In order to represent this, we define *exogenous events* as external actions coming from the environment that may change the values of data fluents. If we consider our case study, we have defined three different exogenous events :

- $photoLost(loc)$ indicates that all the pictures taken in location loc have been lost;
- $fireRisk(loc)$ alerts about a fire that is broken out in location loc ;
- $rockSlide(loc)$ alerts about a rock slide collapsed in location loc .

When an exogenous event is detected, it possibly causes an update in the values of some data fluent. This means that interested fluents has to capture the possible exogenous event modification through their successor state axiom.

Example 3.3.8. The task *takephoto* is used for instructing a service (specifically, a human actor) *act* to move in a location *loc* in order to take some pictures. When *act* completes the execution of the task *takephoto* that was assigned to her/him, the physical outcome q_j of the task is stored in a specific data fluent $PhotoTaken_\varphi$. Note that if an exogenous event *photoLost* is captured in *s*, the value of the fluent $PhotoTaken_\varphi$ (that holds in *s*) switches to false in the situation $do(a, s)$.

$$\begin{aligned}
PhotoTaken_\varphi(loc, do(a, s)) = true &\equiv \\
&(\exists q_j \text{ s.t. } a = \text{release}(act, id, \text{takephoto}, [loc], [true], [q_j]) \wedge (q_j = true)) \vee \\
&(PhotoTaken_\varphi(loc, s) = true \wedge \\
&((a \neq \text{photoLost}(loc)) \vee \\
&\neg \exists q'_j \text{ s.t. } (a = \text{release}(act, id, \text{takephoto}, [loc], [true], [q'_j]) \wedge (q'_j = false))).
\end{aligned} \tag{3.15}$$

The list of all predicates, fluents, exogenous events and abbreviations defined through IndiGolog can be named as the **SitCalc Theory**. Remember that, given a SitCalc Theory, before to execute a process linked to the SitCalc Theory through the IndiGolog PMS, there is the need to clearly specify the values of each data fluent in the initial situation S_0 , that represents the situation in which no actions have yet occurred. In the following example, we show how we set the values of data fluents in S_0 for representing the starting state of our contextual scenario.

Example 3.3.9. We complete the formalization of the scenario described in our case study by providing the whole list of data fluents (all the corresponding successor state axioms are shown in the Appendix), workitems and precondition axioms:

- $At_\varphi(c, s) = loc$ records the location *loc* in which the human service *c* is situated in situation *s*. In the initial situation S_0 , each human service is located in *loc00*.
- $AtRobot_\varphi(c, s) = loc$ records the location *loc* in which the robot service *c* is situated in situation *s*. In the initial situation S_0 , each robot service is located in *loc00*.
- $Evacuated_\varphi(l, s)$ is true if the location *l* is proven to be evacuated in situation *s*. In S_0 , no location has been already evacuated.
- $Status_\varphi(l, s) = st$ records the status of a location *l* in situation *s*, that can be equal to 'ok' (meaning that nothing wrong still happened), or to 'fire' (meaning that a fire has broken out in location *l*) or to 'debris' (meaning that a rock slide has collapsed in *l*). Asynchronously, in any moment, an exogenous event *fireRisk*(*l*) can change the value of $Status_\varphi(l, s)$ to 'fire'. At the same way, an exogenous event *rockSlide*(*l*) can turn the value of $Status_\varphi(l, s)$ to 'debris'. In S_0 , $\forall loc \text{ s.t. } Location_type(loc) \Rightarrow Status_\varphi(l, S_0) = ok$.
- $PhotoTaken_\varphi(l, s)$ is true if in situation *s* some pictures have been taken in location *l*. If an exogenous event *photoLost*(*l*) is captured, the value of $PhotoTaken_\varphi(l, s)$ is turned asynchronously to *false*. In S_0 , no picture has been captured in any location *l*.

- $BatteryLevel_{\varphi}(c, s)$ stores the battery charge level of each robot c in situation s . The level of the robot battery may decrease of a fixed quantity (equal to the value of $MoveStep_{\varphi}(s)$ or to $DebrisStep_{\varphi}(s)$) after c completes a move task or a removedebris task. A human service can recharge the battery of a robot by increasing its level of a fixed quantity equal to the value of $BatteryRecharging_{\varphi}(s)$. For each robot service c , $BatteryLevel_{\varphi}(c, S_0) = 3$.
- $GeneralBattery_{\varphi}(s)$ reflects the total amount of battery contained in the power pack used for recharging the battery of each robot. Each recharging action, that corresponds to the execution of the chargeBattery task, decreases the value of $GeneralBattery_{\varphi}(s)$ of a fixed quantity equal to the value of $BatteryRecharging_{\varphi}(s)$. In the initial situation S_0 , $GeneralBattery_{\varphi}(S_0) = 30$.
- $BatteryRecharging_{\varphi}(s)$ is a data fluent that reflects the amount of each recharging action. In the initial situation S_0 , $BatteryRecharging_{\varphi}(S_0) = 10$.
- $MoveStep_{\varphi}(s)$ indicates the fixed quantity of battery that will be consumed by a robot after the execution of a move task. In the initial situation S_0 , $MoveStep_{\varphi}(S_0) = 2$.
- $DebrisStep_{\varphi}(s)$ indicates the fixed quantity of battery that will be consumed by a robot after the execution of a removedebris task. In the initial situation S_0 , $DebrisStep_{\varphi}(S_0) = 3$.
- $Neigh_{\varphi}(loc1, loc2, s)$ expresses the neighborhood property between two locations $loc1$ and $loc2$. In S_0 , the fluent holds for all that locations considered as adjacent in the scenario depicted in our case study. Hence, $Neigh_{\varphi}(loc00, loc01, S_0) = true$, $Neigh_{\varphi}(loc00, loc10, S_0) = true$, etc.
- $Covered_{\varphi}(loc, s)$ indicates all that locations that are directly covered by the radio range provided by the main antenna. In S_0 , $Covered_{\varphi}(loc00, S_0) = true$, $Covered_{\varphi}(loc01, S_0) = true$, etc.

Now, we show the complete list of the admissible workitems executable by the IndiGolog PMS and precondition axioms for each task.

- The task go can be executed by the PMS if and only if it is associated to a couple of inputs, that represent two valid $Location_type$, and to an expected output, that is again a $Location_type$.

$$\forall (from, to, id) \text{ s.t. } Identifier(id) \wedge Location_type(from) \wedge Location_type(to) \Rightarrow workitem(go, id, [from, to], [to]).$$

A task go with inputs $from$ and to can be assigned to a service act in situation s if and only if act is located in $from$ and is connected to the network.

$$Poss(assign(act, id, go, [from, to], [to]), s) \Leftrightarrow workitem(go, id, [from, to], [to]) \wedge At_{\varphi}(act, s) = from \wedge isConnected(act, s).$$

- The task *move* can be executed by the PMS if and only if it is associated to a couple of inputs, that represent two valid *Location_type*, and to an expected output, that is again a *Location_type*.

$$\forall (from, to, id) \text{ s.t. } Identifier(id) \wedge Location_type(from) \wedge Location_type(to) \Rightarrow workitem(move, id, [from, to], [to]).$$

A task *move* with inputs *from* and *to* can be assigned to a service *rb* in situation *s* if and only if *rb* is located in *from*, is connected to the network and has enough battery charge level for executing the task.

$$Poss(assign(rb, id, move, [from, to], [to]), s) \Leftrightarrow workitem(move, id, [from, to], [to]) \wedge AtRobot_{\varphi}(rb, s) = from \wedge isRobotConnected(rb, s) \wedge BatteryLevel_{\varphi}(rb) \geq MoveStep_{\varphi}(s).$$

- The task *takephoto* can be executed by the PMS if and only if it is associated to an input that represents the *Location_type* where the pictures have to be taken, and to an expected output corresponding to a boolean value indicating the correct execution of the task.

$$\forall (loc, id) \text{ s.t. } Identifier(id) \wedge Location_type(loc) \Rightarrow workitem(takephoto, id, [loc], [true]).$$

A task *takephoto* with input *loc* can be assigned to a service *act* in situation *s* if and only if *act* is located in *loc* and is connected to the network.

$$Poss(assign(act, id, takephoto, [loc], [true]), s) \Leftrightarrow workitem(takephoto, id, [loc], [true]) \wedge At_{\varphi}(act, s) = loc \wedge isConnected(act, s).$$

- The task *evacuate* can be executed by the PMS if and only if it is associated to an input that represents the *Location_type* where it is required to evacuate some people, and to an expected output corresponding to a boolean value indicating the correct execution of the task.

$$\forall (loc, id) \text{ s.t. } Identifier(id) \wedge Location_type(loc) \Rightarrow workitem(evacuate, id, [loc], [true]).$$

A task *evacuate* with input *loc* can be assigned to a service *act* in situation *s* if and only if *act* is located in *loc* and is connected to the network. Moreover, it is required that *loc* is not affected by debris or fire in situation *s*, and it has not already been evacuated.

$$Poss(assign(act, id, evacuate, [loc], [true]), s) \Leftrightarrow workitem(evacuate, id, [loc], [true]) \wedge At_{\varphi}(act, s) = loc \wedge Evacuated(loc, s) = false \wedge Status(loc, s) = ok \wedge isConnected(act, s).$$

- The task `updatestatus` can be executed by the PMS if and only if it is associated to an input that represents the `Location_type` where it is required to update the status of the emergency, and to an expected output indicating that the final status of the location is good (i.e., value 'ok' for the expected output).

$$\forall (loc, st, id) \text{ s.t. } Identifier(id) \wedge Location_type(loc) \wedge Status_type(st) \Rightarrow workitem(updatestatus, id, [loc], [ok]).$$

A task `updatestatus` with input `loc` can be assigned to a service act in situation s if and only if act is located in `loc` and is connected to the network. Moreover, it is required that `loc` is not affected by debris or fire in situation s .

$$Poss(assign(act, id, updatestatus, [loc], [ok]), s) \Leftrightarrow workitem(updatestatus, id, [loc], [ok]) \wedge At_{\varphi}(act, s) = loc \wedge Status_{\varphi}(loc, s) = ok \wedge isConnected(act, s).$$

- The task `extinguishfire` can be executed by the PMS if and only if it is associated to an input that represents the `Location_type` where it is required to extinguish a fire, and to an expected output indicating that after the execution of the task, the fire will be extinguished (i.e., value 'ok' for the expected output).

$$\forall (loc, st, id) \text{ s.t. } Identifier(id) \wedge Location_type(loc) \wedge Status_type(st) \Rightarrow workitem(extinguishfire, id, [loc], [ok]).$$

A task `extinguishfire` with input `loc` can be assigned to a service act in situation s if and only if act is located in `loc` and is connected to the network. Moreover, it is required that in location `loc` a fire broke out in situation s .

$$Poss(assign(act, id, extinguishfire, [loc], [ok]), s) \Leftrightarrow workitem(extinguishfire, id, [loc], [ok]) \wedge At_{\varphi}(act, s) = loc \wedge Status_{\varphi}(loc, s) = fire \wedge isConnected(act, s).$$

- The task `removedebris` can be executed by the PMS if and only if it is associated to an input that represents the `Location_type` where it is required to remove some debris due to a rock slide, and to an expected output indicating that after the execution of the task the debris will be removed (i.e., value 'ok' for the expected output).

$$\forall (loc, st, id) \text{ s.t. } Identifier(id) \wedge Location_type(loc) \wedge Status_type(st) \Rightarrow workitem(removedebris, id, [loc], [ok]).$$

A task `removedebris` with input `loc` can be assigned to a service `rb` in situation s if and only if `rb` is located in `loc`, is connected to the network and has enough battery charge level for executing the task. Moreover, it is required that in location `loc` a rock slide has collapsed in situation s .

$$Poss(assign(rb, id, removedebris, [loc], [ok]), s) \Leftrightarrow workitem(removedebris, id, [loc], [ok]) \wedge AtRobot_{\varphi}(rb, s) = loc \wedge Status_{\varphi}(loc, s) = debris \wedge BatteryLevel_{\varphi}(rb) \geq DebrisStep_{\varphi} \wedge isConnected(act, s).$$

- The task `chargeBattery` can be executed by the PMS if and only if it is associated to an input that represents the robot that needs to be charged. The task has no expected outputs, meaning that we consider a recharging activity as if it always succeeds.

$$\forall (rb, id) \text{ s.t. } Identifier(id) \wedge Service(rb) \Rightarrow \\ workitem(chargeBattery, id, [rb], []).$$

A task `chargeBattery` with input `rb` can be assigned to a service act in situation s if act is connected to the network, `rb` provides a battery (meaning that it is effectively a robot), and both `rb` and act are located in `loc`.

$$Poss(assign(act, id, chargeBattery, [rb], []), s) \Leftrightarrow \\ workitem(chargeBattery, id, [rb], [ok]) \wedge \\ At_{\varphi}(act, s) = AtRobot_{\varphi}(rb, s) \wedge \\ Provides(rb, battery) \wedge isConnected(act, s).$$

3.3.1 Realizing the Framework

In Figure 3.3, we show how `SmartPM` has been concretely coded by the interpreter of `IndiGolog`. The main procedure of the `IndiGolog` program is `Main()`, which involves three interrupts running at different priorities. Each interrupt is guarded by the fluent `Finished`; if it holds, it means that the process execution has been completed successfully.

$$Finished(do(a, s)) = true \equiv \\ a = finish \vee (Finished(s) = true). \quad (3.16)$$

The procedure `Monitor()`, which runs at higher priority, is in charge of monitoring changes in the environment and adapting accordingly. The first step in `Monitor` checks whether fluent `RealityChanged` holds true, meaning that a task has completed its execution or that an exogenous (unexpected) action has occurred in the system.

$$RealityChanged(do(a, s)) = true \equiv \\ \exists t, c, id, \vec{i}, \vec{p}, \vec{q} \text{ s.t. } a = release(c, id, t, \vec{i}, \vec{p}, \vec{q}) \vee \\ exogenous(a) \vee \\ (RealityChanged(s) = true \wedge a \neq resetReality). \quad (3.17)$$

If some change in the contextual data is considered as *Relevant*, the procedure `Adapt()` is launched for the synthesis of the recovery procedure. Both if *Relevant* holds or not, the `Monitor()` concludes by executing the action `resetReality`, which turns the fluent `RealityChanged` to *false*. We give more details about the *Relevant* fluent and the working of the `Monitor()` procedure in the following sections.

At a lower priority, the system runs the actual `IndiGolog` program representing the dynamic process to be executed, namely procedure `Process()`. As shown in Fig. 3.3, the process is composed by three branches of tasks to be executed in parallel (cf. Section 3.2.2 for the list of `IndiGolog` constructs), that correspond exactly to the emergency management process defined in Section 1.3. For example, the procedure `Branch1()` depicts three tasks to be executed in sequence. Specifically, `Branch1()`

```

PROC Main()
1   $\langle \neg Finished \wedge RealityChanged \rightarrow [Monitor()] \rangle$ ;
2   $\langle \neg Finished \rightarrow [Process(); finish] \rangle$ ;
3   $\langle \neg Finished \rightarrow [wait] \rangle$ ;

PROC Monitor()
1  if Relevant
2    then Adapt();
3  resetReality;

PROC ManageExecution(Task, Id, Inputs, ExpectedOutputs)
1   $\pi(Srv)$ ; ( $Capable(Srv, Task) \wedge Free(Srv)$ )?
2  assign(Srv, Id, Task, Inputs, ExpectedOutputs).
3  start(Srv, Id, Task).
4  ackCompl(Srv, Id, Task).
5  release(Srv, Id, Task, Inputs, ExpectedOutputs, PhysicalOutputs).

PROC Process()
1  [Branch1()||Branch2()||Branch3()].

PROC Branch1()
1  [ManageExecution(go, id1, [loc00, loc33], [loc33]);
2  ManageExecution(takephoto, id2, [loc33], [true]);
3  ManageExecution(updatestatus, id3, [loc33], [ok])].

PROC Branch2()
1  [ManageExecution(go, id4, [loc00, loc32], [loc32]);
2  ManageExecution(evacuate, id5, [loc32], [true]);
3  ManageExecution(updatestatus, id6, [loc32], [ok])].

PROC Branch3()
1  [ManageExecution(go, id7, [loc00, loc31], [loc31]);
2  ManageExecution(evacuate, id8, [loc31], [true]);
3  ManageExecution(updatestatus, id9, [loc31], [ok])].

```

Figure 3.3. The core procedures of SmartPM.

instructs first a selected service to reach location *loc33* starting from location *loc00*, then to take pictures in that location (the expected output of this task is *true*, meaning that the pictures have been correctly collected), and finally to update the status of location *loc33* (the 'ok' value means that nothing strange, such as a fire or a rock slide, has happened in *loc33*). Note that in Fig. 3.3 we make use of a Prolog-like notation and lists of inputs/expected outputs/physical outputs are enclosed between squared brackets.

Each **Branch()** procedure relies, in turn, on procedure **ManageExecution**. Given a specific task *T*, the procedure **ManageExecution** first searches for a service *c* that is *Free* in situation *s* and is *Capable* to execute the task and, if such a service exists, the procedure will be in charge to manage task assignment, start signaling, acknowledgment of completion, and final release.

In the list of interrupts, at the lowest priority, there is also a third possibility for managing the progression of the process, that is activated only if the process

is still not finished, but for some reason, it can not progress in its execution. The third interrupt, that consists just in waiting, reflects the fact that, for example, a certain task can not be assigned to any qualified service (i.e., the π function is unable to find any service providing all the capabilities required by that task) or all qualified services are currently involved in the performance of other tasks. If the third interrupt is activated, the control passes back to the process designer, which can manually manage the situation (for example, by adding new services or by updating the capabilities of the existing services).

3.4 Monitoring for Failures

In this section, we turn our attention to the mechanism for automatically detecting failures. As described in Section 3.1, adaptation in SmartPM can be seen as reducing the gap between the *expected reality*, the (idealized) model of reality that is used by the PMS to reason, and the *physical reality*, the real world with the actual values of conditions and outcomes. In order to understand how the proposed technique works, we start by formalizing the concepts of *physical* and *expected reality*.

Definition 1. A *physical reality* $\Phi(s)$ is the set of all data fluents $X_{\varphi,y}$ (where y ranges over $\{1..m\}$) defined in the SitCalc theory. Hence, $\Phi(s) = \bigcup_{y=1..m} \{X_{\varphi,y}\}$.

The physical reality $\Phi(s)$ captures exactly the value assumed by each *data fluent* in the situation s . Such value reflects what is really happening in the real environment whilst the process is under execution. However, the PMS must guarantee that each task in the process is executed correctly, i.e., with an output that satisfies the process specification. For this purpose, the concept of *expected reality* $\Psi(s)$ is needed. For some data fluents, the ones affected by a *supposed* task effect (cf. Section 3.3), we introduce a new *expected fluent* X_{ψ} that is meant to record the “expected” value of X after the execution of a task T . The successor state axiom for this new fluent is straightforward:

$$\begin{aligned}
X_{\psi}(\vec{\sigma}, do(a, s)) = p_j &\equiv \\
&(\exists c, id, \vec{i}, \vec{p}, \vec{q} \text{ s.t. } a = \text{release}(c, id, T, \vec{i}, \vec{p}, \vec{q}) \wedge \\
&\vec{p} = [p_0, \dots, p_j, \dots, p_w]) \vee \\
&(X_{\psi}(\vec{\sigma}, s) = p_j \wedge \neg \exists c, id, \vec{i}, \vec{p}', \vec{q} \text{ s.t. } a = \text{release}(c, id, T, \vec{i}, \vec{p}', \vec{q}) \wedge \\
&\vec{p}' = [p'_0, \dots, p'_j, \dots, p'_w] \wedge (p'_j \neq p_j)).
\end{aligned} \tag{3.18}$$

The fluent states that, in the expected reality, a task is *always* executed correctly (also when it is not), by forcing the value of X_{ψ} to the value of the expected output p_j .

Example 3.4.1. When an actor *act* terminates the execution of the task *go* from a location *from* to a location *to*, the expected fluent At_{ψ} will assume the value *to* without considering the real outcome of the task.

$$\begin{aligned}
At_{\psi}(act, do(a, s)) = to &\equiv \\
&(\exists q_j \text{ s.t. } a = \text{release}(act, id, go, [from, to], [to], [q_j])) \vee \\
&(At_{\psi}(act, s) = to \wedge \\
&\neg \exists to', q_j \text{ s.t. } a = \text{release}(act, id, go, [from, to'], [to'], [q_j]) \wedge (to' \neq to)).
\end{aligned} \tag{3.19}$$

Therefore, an expected fluent X_ψ holds every time a task is completed with any real effect, by storing one of the expected outcomes of the task. More precisely, given a task T whose execution returns a list of real outcomes $\vec{q} = [q_0, \dots, q_j, \dots, q_w]$ and a list of expected outcomes $\vec{p} = [p_0, \dots, p_j, \dots, p_w]$, there exists a data fluent X_φ and a corresponding expected fluent X_ψ that will assume respectively the values q_j and p_j . We need to underline that some data fluents do not provide an associated expected fluent; for example, the fluent $BatteryLevel_\varphi$ records the battery charge level of each robot, and this value changes of a fixed quantity after each robot movement, whatever is the real final position of the robot. On the contrary, given a data fluent X_φ that records one of the *supposed* effects of a task, an expected fluent X_ψ could be defined. For example, although the task *move* provides a “supposed” outcome, there is no need to record the expected position of a robot. In fact, if a robot reaches a location different from the one expected, the important aspect is that this physical position allows in any way to guarantee the network connection.

Definition 2. *An expected reality $\Psi(s)$ is the set of all expected fluents $X_{\psi,y}$ (where y ranges over $\{1..k\}$, with $k \leq m$) defined in the SitCalc theory. Hence, $\Psi(s) = \bigcup_{y=1..k} \{X_{\psi,y}\}$.*

In Figure 3.4, the overall framework for detecting failures is depicted. The PMS starts by taking in input a process specification δ_0 formalized in IndiGolog, and builds the SitCalc theory (i.e., the set of first-order predicates and situation calculus fluents) representing the contextual environment in which the process has to be executed. Finally, the PMS instantiates the initial situation S_0 , that indicates starting values for each data and expected fluent.

The execution of a process can be interrupted by the *Monitor* module when a misalignment between the physical and the expected reality is discovered. Specifically, the Monitor blocks the execution of the main process δ (i) after the occurrence of a *release* action in the current situation s , meaning that some service has completed the execution of a task, or (ii) if some exogenous event has been caught in s . While in the first case there is the possibility that some physical and expected fluent has changed its value, an exogenous event may affect directly only some data fluent (i.e., the physical reality).

In this sense, our framework is able to capture - and to recover from - two different kinds of task failure. An *internal failure* is related to the failure in the execution of a task, i.e., the task does not terminate, or it is completed with an output that differs from the expected one. An *external failure* is represented as an exogenous event e , given in input by the external environment, that forces a set of data fluents to assume a value imposed by the event itself. Such a new value could differ from the expected one, by generating a discrepancy between the two realities.

A recovery procedure is needed if the two realities are different from each other, i.e., some tasks in the process failed their execution by returning an output q whose value is different from the expected output p , or if some exogenous event has modified the physical reality in a undesirable way.

The reduction of this gap requires sufficient knowledge of both kinds of realities. This knowledge, recorded in data and expected fluents, allows the PMS to sense deviations and to modify the process to ensure that, at the end, the above gap will




Figure 3.4. Execution monitoring in SmartPM.

be removed. Formally, a situation s is known as *Relevant*- candidate for adaptation - iff :

$$Relevant(\delta, s) \equiv \neg SameState(\Phi(s), \Psi(s)) \quad (3.20)$$

Predicate $SameState(\Phi(s), \Psi(s))$ holds iff the states denoted by $\Phi(s)$ and $\Psi(s)$ are the same².

Definition 3. Given a situation s and a set \vec{F} of fluents , a state($\vec{F}(s)$) is the set composed by the values - in situation s - of each fluent F_j that belongs to \vec{F} . Hence, $state(\vec{F}(s)) = \bigcup_{j=1..m} \{F_j\}$ s.t. $F_j \in \vec{F}$.

Next, let us formalize how the monitor works. After the termination of a task, or after the occurrence of an exogenous event, the PMS infers the new situation s' derived from the execution of a *release* action (or from the arrival of an exogenous event) and passes it to the Monitor, together with the fragment of process δ' remaining to be executed. In case of task completion in situation s , the realities $\Phi(s)$ and $\Psi(s)$ will evolve in $\Phi(s')$ and $\Psi(s')$, while in case of an exogenous event, only $\Phi(s)$ evolves in $\Phi(s')$, while $\Psi(s')$ remains equal to $\Psi(s)$. Finally, it is obvious that in case of exogenous events, $\delta' = \delta$ (i.e., an exogenous event does not reflect a task completion, but a change in the physical reality).

Now the Monitor component checks if $Relevant(\delta', s')$ holds. If it does not hold, it means that no adaptation is required for the process δ' , that can be still executed

²The evaluation of $SameState$ is performed only for those data fluents for which there exists a corresponding expected fluent. It is obvious that for data fluents considered as *not relevant* for adaptation there is no need to monitor their evolution over situations.

in situation s' . In such a case, the Monitor put δ'' equal to δ' and carries on with the main process execution.

On the contrary, if $Relevant(\delta', s')$ holds, adaptation of the process δ' is needed. Specifically, the purpose of SmartPM is to devise a recovery procedure δ_a that turns the *wrong* physical reality $\Phi(s')$ in the *correct* expected reality $\Psi(s')$. The SmartPM system proposes two different adaptation mechanisms for recovering the process from a failure or for adapting it to changing circumstances in the contextual environment due to an exogenous event. The first technique consists in using the lookahead search construct Σ provided by IndiGolog, while the second technique, which is the one currently deployed on the SmartPM prototype, is based on classical planning algorithms. Both the adaptation techniques are able to synthesize a linear process δ_a , i.e., a process consisting of a sequence of tasks, such that $\delta_a = [t_{a_0}; \dots; t_{a_n}]$. Such a recovery procedure δ_a will be inserted just before δ' , by devising a new process $\delta'' = (\delta_a; \delta')$. Note that whenever a process needs to be adapted, every running task is interrupted, since the “repair” sequence of tasks δ_a is placed before them. Thus, active branches can only resume their execution after the repair sequence δ_a has been executed. This last requirement is fundamental to avoid the risk of introducing data inconsistencies during the repair phase.

The adapted process $\delta'' = (\delta_a; \delta')$ must guarantee some properties in order to ensure the correctness of the recovery mechanism:

- The recovery process δ_a has to be synthesized for being executed in situation s' . Moreover, after the execution of δ_a , it should be guaranteed that $\Phi(s')$ is turned in $\Psi(s')$. We will discuss this aspect in the following sections.
- The execution of the recovery procedure δ_a corresponds to a new situation s'' and to new realities $\Phi(s'')$ and $\Psi(s'')$. In order to ensure that the remaining part of the main process δ' is still executable in s'' , the recovery mechanism should guarantee that $\Psi(s')$ and $\Psi(s'')$ represent the same expected state of the world, i.e., the fluent $SameState(\Psi(s'), \Psi(s''))$ must hold after the execution of δ_a . Specifically, we want that after the execution of the recovery process δ_a in situation s' (that results in a new situation s''), the remaining process δ' to be executed in s'' is equivalent to execute δ' in s' .

We formalize this second point by exploiting the concept of *bisimulation*. In Computer Science, two systems are *bisimilar* if they match each other’s moves, i.e., one system simulates the other and vice-versa. In this sense, each of the systems can not be distinguished from the other by external an observer. For proving this, we define the the predicate $SameConfig(\delta_\alpha, s_\alpha, \delta_\beta, s_\beta)$ as follows:

Definition 1. A predicate $SameConfig(\delta_\alpha, s_\alpha, \delta_\beta, s_\beta)$ is correct if for every $\delta_\alpha, s_\alpha, \delta_\beta, s_\beta$:

1. $Final(\delta_\alpha, s_\alpha) \Leftrightarrow Final(\delta_\beta, s_\beta)$
2. $\forall a, \delta_\alpha$ s.t. $Trans(\delta_\alpha, s_\alpha, \overline{\delta_\alpha}, do(a, s_\alpha)) \Rightarrow \exists \overline{\delta_\beta}$ s.t. $Trans(\delta_\beta, s_\beta, \overline{\delta_\beta}, do(a, s_\beta)) \wedge SameConfig(\overline{\delta_\alpha}, do(a, s_\alpha), \overline{\delta_\beta}, do(a, s_\beta))$
3. $\forall a, \delta_\beta$ s.t. $Trans(\delta_\beta, s_\beta, \overline{\delta_\beta}, do(a, s_\beta)) \Rightarrow \exists \overline{\delta_\alpha}$ s.t. $Trans(\delta_\alpha, s_\alpha, \overline{\delta_\alpha}, do(a, s_\alpha)) \wedge SameConfig(\overline{\delta_\beta}, do(a, s_\beta), \overline{\delta_\alpha}, do(a, s_\alpha))$

Intuitively, a predicate $SameConfig(\delta_\alpha, s_\alpha, s_\beta, \delta_\beta)$ is said to be correct if δ_α and δ_β are terminable either both or none of them. Furthermore, for each action a performable by δ_α in the situation s_α , there exists the same action a performable by δ_β in the situation s_β (and viceversa). Moreover, the resulting configurations $(\overline{\delta_\alpha}, do(a, s_\alpha))$ and $(\overline{\delta_\beta}, do(a, s_\beta))$ must still satisfy $SameConfig$.

In our case, we can adopt a specific definition for $SameConfig$, that we call $SameConfig_{PM}(\delta_\alpha, s_\alpha, s_\beta, \delta_\beta)$:

$$\begin{aligned} SameConfig_{PM}(\delta_\alpha, s_\alpha, s_\beta, \delta_\beta) \Leftrightarrow \\ SameState(\Psi(s_\alpha), \Psi(s_\beta)) \wedge \delta_\alpha = \delta_\beta \end{aligned} \quad (3.21)$$

In other words, $SameConfig_{PM}$ states that δ_α, s_α and δ_β, s_β are the same configuration if (i) all fluents have the same truth values in both $\Psi(s_\alpha)$ and $\Psi(s_\beta)$ ($SameState$), and (ii) δ_β is equal to δ_α (i.e., both δ_α and δ_β correspond to the remaining process to be executed δ'). The following shows that $SameConfig_{PM}$ is correct.

Theorem 1. $SameConfig_{PM}(\delta_\alpha, s_\alpha, s_\beta, \delta_\beta)$ is correct.

Proof. We show that $SameConfig_{PM}$ is a bisimulation. Indeed:

- Since $SameState(\Psi(s_\alpha), \Psi(s_\beta))$ requires all expected fluents to have the same values both in s_α and s_β , we have that $(Final(\delta_\alpha, s_\alpha) \Leftrightarrow Final(\delta_\beta, s_\beta))$.
- Since $SameState(\Psi(s_\alpha), \Psi(s_\beta))$ requires all expected fluents to have the same values both in s_α and s_β , it follows that the PMS is allowed for the same process δ_α to assign the same tasks both in s_α and in s_β and moreover for each action a and situation s_α and s_β s.t. $SameState(\Psi(s_\alpha), \Psi(s_\beta))$, we have that $SameState(\Psi(do(a, s_\alpha)), \Psi(do(a, s_\beta)))$ hold. As a result, for each a and $\overline{\delta_\alpha}$ such that $Trans(\delta_\alpha, s_\alpha, \overline{\delta_\alpha}, do(a, s_\alpha))$ we have that $Trans(\delta_\alpha, s_\beta, \overline{\delta_\alpha}, do(a, s_\beta))$ and $SameConfig_{PM}(\overline{\delta_\alpha}, do(a, s_\alpha), \overline{\delta_\beta}, do(a, s_\beta))$. Similarly for the other direction.

□

We can now formalize a predicate $Recovery$ that describes how our recovery mechanism works:

$$\begin{aligned} Recovery(\delta', s', s'', \delta'') \Leftrightarrow \\ \exists \delta_a \text{ s.t. } \delta'' = \delta_a; \delta' \wedge Linear(\delta_a) \wedge \\ Do(\delta_a, s', s'') \wedge \neg Relevant(\delta', s'') \wedge \\ SameConfig_{PM}(\delta', s', s'', \delta') \end{aligned} \quad (3.22)$$

Intuitively, $Recovery(\delta', s', s'', \delta'')$ holds if the program δ' that was intended to be executed in s' is adapted in a new program δ'' . The adapted process δ'' is a sequence composed by the linear recovery procedure δ_a and the remaining part of the process to be executed δ' . The execution of δ_a in s' results in a new situation s'' , where $SameState(\Phi(s''), \Psi(s''))$ (i.e., $\neg Relevant(\delta', s'')$) and $SameConfig_{PM}(\delta', s', s'', \delta')$ hold.

The nice feature of $Recovery$ is that it searches for a linear program that achieves a certain formula, namely $SameState(\Phi(s), \Psi(s))$. Moreover, restricting to sequential

programs obtained by planning with no concurrency does not prevent any recoverable process from being adapted. In sum, we have reduced the synthesis of a recovery program to a classical Planning problem in AI [90]. As a result we can adopt a well-developed literature about planning for our aim. In particular, if services and input/output parameters are finite, then the recovery can be reduced to *propositional* planning, which is known to be decidable in general.

3.5 The SmartPM Adaptation Mechanisms

3.5.1 The Built-in Adaptation Mechanism

The first adaptation technique we analyze for SmartPM is based on synthesizing a number of candidate recovery processes and on simulating them off-line. More precisely, when the fluent *Relevant* holds, we need to find a recovery procedure which is able to align the physical reality $\Phi(s)$ with the expected reality $\Psi(s)$. A quick solution to this problem consists of devising the recovery plan on-line, by trying to execute performable actions in $\Phi(s)$ and senses what the next action should be, on the basis of a “distance” notion between the two realities. Such a solution does not require a reasoner to determine a lengthy course of action (formed perhaps of hundreds of tasks) before executing the first step in the world.

However, on the other hand, once an action has been executed in the world, there may be no way of backtracking it if it is later found out that it was performed incorrectly. This aspects assumes a great value if we think that we are executing tasks in an emergency management context, where delays or useless/incorrect activities may easily prevent the correctness of the whole procedure. As a result, an on-line execution of a program may fail where an off-line execution would succeed.

On Fig. 3.6, we show the fragment of the IndiGolog code related to our built-in adaptation mechanism. Procedure **Adapt()** starts by invoking a basic action *adaptStart*, whose effect is to make the fluent *Adapting* equal to *true*.

$$\begin{aligned} \text{Adapting}(\text{do}(a, s)) = \text{true} \equiv \\ a = \text{adaptStart} \vee \\ (\text{Adapting}(s) = \text{true} \wedge a \neq \text{adaptFinish}). \end{aligned} \quad (3.23)$$

Then, **Adapt()** invokes the procedure **AdaptingProgram()** in order to build and execute the recovery program and, at the same time (cf. the prioritized concurrency in line 2), it waits until the recovery procedure has been completely performed. This will happen when **AdaptingProgram()** will terminate its execution, and the basic action *adaptFinish* will turn the fluent *Adapting* to *false*.

Procedure **AdaptingProgram()** relies on the IndiGolog *search operator* Σ , that provides a form of lookahead planning. The idea is that given any program δ , the program $\Sigma(\delta)$ executes online just like δ does off-line. In other words, before taking any action, it first ensures using online reasoning that this step can be followed successfully by the rest of δ . More precisely, according to [26], the semantics of the search operator is that:

$$\text{Trans}(\Sigma(\delta), s, \Sigma(\delta'), s') \Leftrightarrow \text{Trans}(\delta, s, \delta', s') \wedge \exists s^* \text{ s.t. } \text{Do}(\delta', s', s^*).$$


```

PROC Adapt()
1  adaptStart.
2  (AdaptingProgram(); adaptFinish) » (while Adapting do wait).

PROC AdaptingProgram()
1   $\Sigma$ ([SearchProgram()],
2  assumptions[{assign(Srvc, Id, Task, Inputs, ExpectedOutputs)}, readyToStart(Srvc, Id, Task)],
3  {start(Srvc, Id, Task), finishedTask(Srvc, Id, Task, ExpectedOutputs)}]).

PROC SearchProgram()
1  Plans(1, 10).

PROC Plans(m, n)
1  (m ≤ n)?
2  [ActionSequence(m); (-Relevant)?] | Plans(m + 1, n).

PROC ActionSequence(m)
1  if m > 0
2  then  $\pi$ (Task, Id, Inputs, ExpectedOutputs); workitem(Task, Id, Inputs, ExpectedOutputs)?
3      ManageExecution(Task, Id, Inputs, ExpectedOutputs).
4      ActionSequence(m - 1).

PROC ManageExecution(Task, Id, Inputs, ExpectedOutputs)
1   $\pi$ (Srvc); (Capable(Srvc, Task) ∧ Free(Srvc))?
2  assign(Srvc, Id, Task, Inputs, ExpectedOutputs).
3  start(Srvc, Id, Task).
4  ackCompl(Srvc, Id, Task).
5  release(Srvc, Id, Task, Inputs, ExpectedOutputs, PhysicalOutputs).

```

Figure 3.5. The built-in adaptation mechanism of SmartPM.

If δ is the entire program under consideration, $\Sigma(\delta)$ emulates complete off-line execution. In SmartPM, we use a specialized version of Σ that relies on some assumptions on the performable actions. Each assumption is of form $\{actionPMS(\vec{x}), actionService(\vec{y})\}$, meaning that any action $actionPMS$ executed by the SmartPM engine with input \vec{x} will be eventually “complemented” by action $actionService$ executed by a service with input \vec{y} . Vector of parameters \vec{y} is a fully-deterministic transformation of \vec{x} . Here we are using the simple case where \vec{y} is a subset to \vec{x} , but one can customize for specific tasks/actions.

Specifically, for SmartPM, we have coded two assumptions: the first is that the action $readyToStart(Srvc, Id, Task)$ performed by a certain service $Srvc$ is expected to follow the PMS action $assign(Srvc, Id, Task, Inputs, ExpectedOutputs)$; the second concerns the PMS action $start(Task, Id, Srvc)$, which is supposed to come before the action $finishedTask(Srvc, Id, Task, Inputs, PhysicalOutputs, ExpectedOutputs)$, executed by $Srvc$. (cf. also Fig. 3.2)

Finally, let us focus on the actual program in charge of building the recovery plan, namely procedure **SearchProgram()** that, in turn, invokes procedure **Plans(m, n)**. Generally speaking, such procedure will try to reach a situation in which *Relevant* does not hold anymore (see line 2). For this aim, the **Plans** procedure builds recovery plans of growing length and simulates their execution starting from $\Phi(s)$. Recovery plans may consist of a variable number of tasks to be executed in sequence, that range from m to n tasks.

Let us suppose, for example, that $m = 1$, meaning that we are executing

Plans(1, 10) for the first time. As a consequence, the procedure **ActionSequence**(1) is invoked. It tries to generate all sequences composed by a single task. Then it simulates their execution in $\Phi(s)$ (cf. the invocation of **ManageExecution** in line 3) and checks if at least one of those sequences can turn the value of $\Phi(s')$ (s' results from the simulated execution of the candidate recovery process) such that $SameState(\Phi(s'), \Psi(s'))$ (cf. line 2 of the procedure **Plans**). If it does not happen, the value of m is increased of one unit (cf. line 2 of the procedure **Plans**) and **Plans**(2, 10) is invoked.

Now, **ActionSequence**(2) searches for sequences of two tasks $[t_1; t_2]$ to be executed in $\Phi(s)$. The execution of each candidate recovery plan happens off-line (i.e., when the main process has been stopped for waiting for the building of the recovery plan itself), but the built-in algorithm simulates an on-line execution of the plan, by applying the effects of the sequence directly on the current $\Phi(s)$. To be more precise, **ActionSequence**(2) first picks an admissible workitem (cf. line 2 of the procedure **ActionSequence**), and then simulates its execution on $\Phi(s)$, by obtaining a new reality $\Phi(s')$. Now, starting from $\Phi(s')$, **ActionSequence** is invoked again for searching any task that is executable on $\Phi(s')$, by devising the new reality $\Phi(s'')$. Again, the control is passed back to **Plans**, that verifies if $SameState(\Phi(s''), \Psi(s''))$.

We need to underline that every expected reality $\Psi(s_\alpha)$ devised in any situation s_α (starting from the situation s) that comes from the simulated execution of recovery processes of growing length, will be always equal to $\Psi(s)$ (i.e., to the expected reality stored after a failure has been sensed). In fact, when the fluent *Adapting* holds, meaning that the system is searching and simulating the execution of some recovery procedure, the system is not allowed to change the values of expected fluents. This because the built-in algorithm tries to execute each task of the recovery procedure as if it returns its expected outcomes, with the purpose to turn $\Phi(s)$ into $\Psi(s)$. This means that $SameState(\Psi(s), \Psi(s_\alpha))$ and, consequently, $SameConfig_{PM}(\delta', s, s_\alpha, \delta')$ hold. Moreover, if the Σ search operator is able to find a recovery procedure δ_a composed by a sequence of tasks such that $Do(\delta_a, s, s_\alpha)$, then also $Recovery(\delta', s, s_\alpha, (\delta_a; \delta'))$ is *true*.

In general, the search technique is iterative deepening: if there exists no sequence whose length is less or equal to m tasks, it tries with length- $(m + 1)$ task sequences. Observe also that **ActionSequence** uses a simple breadth-first search mechanism, which specializes what proposed in [107]. The procedure tries to generate all admissible sequences of tasks composed by m task. This keeps going deeper and deeper till reaching a sequences of 10 tasks or any task sequence that recovers. If no task sequence of at most 10 tasks exists, it is assumed that no recovery is possible. We have adopted 10 as bound to the length of the sequence as it is a reasonable assumption in our scenario.

It is worth highlighting that, because the monitor runs at a higher-priority level than the actual process, the solution plan found for the recovery program $\Sigma[\mathbf{SearchProgram}]$ would run at higher-priority than program **Process**. So, the program **Process** cannot progress until the recovery has been finished and applied. Consequently, after a sensed deviation, the program executed will be equivalent to $(\Sigma[\mathbf{SearchProgram}()]; \delta')$, where δ' is the program remaining from procedure **Process**.

```

PROC Adapt()
1  adaptStart.
2  InvokePlanner().
3  adaptFinish.

PROC InvokePlanner()
1  Init =  $\Phi(s) \cup \{\bigcup_{i=1..k, j=1..n} \{Provides(c_i, b_j)\}\} \cup \{\bigcup_{i=1..k} \{Free(c_i, s)\}\}$ .
2  Goal =  $\Psi(s)$ .
3  send(Planner, [Init, Goal]).
4  receive(Planner, ( $\delta_a; \delta'$ )).
5  alignRealities.

```

Figure 3.6. The plan-based adaptation mechanism of SmartPM.

Theorem 2. *Let assume a domain in which services and input and output parameters are finite. Then given a process δ' and situations s' and s'' , it is decidable to compute a recovery process δ'' with the built-in approach just devised such that $Recovery(\delta', s', s'', \delta'')$ holds.*

Proof. In domains in which services and input and output parameters are finite, also actions and fluents instantiated with all possible parameters are finite. Hence we can phrase the domain as a propositional one and the thesis follows from decidability of propositional planning [90]. \square

3.5.2 The Plan-based Adaptation Approach

We now turn our attention to the plan-based adaptation mechanism currently working in the SmartPM system. Before starting the execution of the process δ , the PMS builds the PDDL representation of each task defined in the SitCalc theory and sends it to an external planner. More in detail, SmartPM is able to build a PDDL planning domain starting from a *domain theory* defined through SmartML, our declarative language used for representing contextual properties of a dynamic scenario. We show the syntax of SmartML in Section 4.3.1, but we anticipate that a SmartML specification can be easily converted in a SitCalc theory (cf. Section 4.3.3).

When a misalignment between $\Phi(s')$ and $\Psi(s')$ is sensed (we consider s' as the situation where something wrong has happened and δ' as the faulty process), the **Monitor()** launches the **InvokePlanner()** procedure for recovering the faulty process δ' . **InvokePlanner** starts by building a PDDL planning problem that reflects the gap between the two realities. Specifically, it first determines the initial state *Init* of the planning problem, by making it equal to $\Phi(s')$. In addition to the values of data fluents, *Init* can also include the values of $Provides(c, b)$ and of $Free(c, s')$, for each capability b and service c defined in SitCalc theory. Such values will be used by the planner for scheduling correctly recovery tasks during the building of the plan.

Then, **InvokePlanner** builds the *goal* of the planning problem. Since our target is to convert $\Phi(s')$ into $\Psi(s')$, it is clear that the *goal* will be equal to $\Psi(s')$. There exists the possibility to add to the *goal* some situation calculus abbreviation that is considered to be relevant for adaptation. For example, in our case study we need to guarantee that each service is connected to the network. Therefore, the *goal* can be

augmented with information stating that the execution of the synthesized recovery plan would not prevent the connection of services to the main network. Since our *goals* are basically a conjunction of literals, the presence of SitCalc abbreviations does not prevent the reachability of the atomic goals contained in $\Psi(s')$ (unless the SitCalc abbreviation directly contradicts some expected fluents value, but this would be considered as a design-time error). The only effect deriving from the presence of SitCalc abbreviations within the *goal* definition is that the searching for a recovery plan will be a bit more constrained.

Example 3.5.1. Let us consider the example described in our case study. Specifically, suppose, for example, that the task $go(loc00, loc33)$ is assigned to actor $act1$ (cf. Fig. 1.3(a)), which reaches instead the location $loc03$ (cf. Fig. 1.3(b)). This means that $act1$ is now located in a different position than the desired one, and s/he is out of the optimal network range. Therefore, after the *release* action has been executed, the fluent At_φ takes the value $loc03$. But this output does not satisfy the expected outcome. The expected output $loc33$ is stored in the fluent At_ψ ; it generates a discrepancy between $\Phi(s')$ and $\Psi(s')$. This means that $Relevant(\delta, s')$ holds, and the main process δ' needs to be adapted.

Now, the IndiGolog engine invokes an external planner by giving as inputs the *init* and *goal* sets just computed. We have to underline that the communication between the IndiGolog engine and the planner is mediated by a **Synchronizer** component, that launches the planner and, when a recovery plan has been computed, translates the plan in a new IndiGolog adapted process $\delta'' = (\delta_a; \delta')$. We give more information about the **Synchronizer** in Chapter 4.

The *receive* command in line 4 of the procedure **InvokePlanner** is blocking, i.e., the PMS waits until the recovery procedure has been devised. If the recovery procedure is empty, it means that no plan exists for the current planning problem, and the control passes back to the process designer, that can try to manage manually the exception. Otherwise, if a recovery process $\delta'' = (\delta_a; \delta')$ is returned, the IndiGolog engine can start to execute it. If we consider the example 3.5.1, the recovery plan needed for removing the gap between the two realities is shown in Fig. 1.4(a).

The last command of the procedure **InvokePlanner** is used for forcing $\Psi(s')$ to the current value of $\Phi(s')$. This because the execution of the recovery plan δ_a will turn the initial state (i.e., $\Phi(s')$) into the goal state (i.e., $\Psi(s')$). Moreover, after the execution of every recovery task of δ_a the situation s' will evolve in a situation s'' , where $\Psi(s'')$ is exactly the goal of the planning problem just reached. This means that $SameConfig_{PM}(\delta', s', s'', \delta')$ holds and, since the execution of the recovery plan guarantees that also $SameState(\Phi(s''), \Psi(s''))$ holds, the predicate $Recovery(\delta', s', s'', (\delta_a; \delta'))$ is *true*.

Theorem 3. *Let assume a domain in which services and input and output parameters are finite. Then, given a process δ' and situations s' and s'' , it is decidable to compute a recovery process δ'' with the planning-based approach just devised such that $Recovery(\delta', s', s'', \delta'')$ holds.*

Proof. The decidability of the plan-based approach relies on the external planner used for the synthesis of the recovery procedure. In the SmartPM system, we have




Figure 3.7. Traditional “plan then execute” cycle (a) and the *continuous planning* approach.

used the LPG-td planner [45] for synthesizing recovery plans. LPG-td is a state-based planner that is based on a stochastic local search in the space of particular “action graphs” derived from the planning problem specification. The basic search scheme of LPG-td is inspired to Walksat [118], an efficient procedure for solving SAT-problems. If a domain in which services and input and output parameters are finite, the LPG-td planner guarantees to terminate [45]. □

3.5.3 The Continuous Planning Approach

In this section, we show a third approach for adapting a process from failures or exogenous events. It is based on *Continuous Planning* techniques and, if compared with the two adaptation approaches employed on SmartPM and discussed above, it provides two interesting features in adapting a process: (i) it is a non-blocking technique, i.e., it does not stop directly any task in the main process during the computation of the recovery procedure and (ii) it allows concurrent branches in the recovery procedure. This approach, developed on the same SitCalc Theory defined in Section 3.3, requires to modify the working of the **Monitor** procedure, and its implementation on the IndiGolog platform is currently on going.

Continuous Planning [145] refers to the process of planning in a world under continual change, where the planning problem is often a matter of adapting to the world when new information is sensed. A continuous planner is designed to persist indefinitely in the environment. Thus it is not a “problem solver” that is given a single goal and then plans and acts until the goal is achieved; rather, it lives through a series of ever-changing goal formulation, planning, and acting phases. Rather than thinking of the planner and execution monitor as separate processes, one of which passes its results to the other, we can think of them as a single process (cf. Fig. 3.7).

The proposed idea is to build the recovery procedure δ_a *in parallel* with the execution of the main process δ , avoiding to stop directly any task in the main process. Once ready, δ_a will be inserted as a new branch of δ and will be executed in concurrency with every other task. This means that the recovery branch will turn δ into a new process $\delta'' = (\delta || \delta_a)$. Before to analyze how the whole approach works, we need to provide some further formal definition.

Since we are working in a domain in which services and input/output parameters are finite, each task defined in δ affects (or is affected by) only a finite number of fluents. This means that each task is interested only in that fragment of reality it contributes to modify.

Definition 4. A task T affects a data/expected fluent X iff $\exists c, id, \vec{i}, \vec{p}, \vec{q}, a$ s.t. $a = \text{release}(c, id, T, \vec{i}, \vec{p}, \vec{q})$. We denote it with $T \triangleright X$.

Definition 5. A task T is affected by a data/expected fluent X iff $\exists c, id, \vec{i}, \vec{p}, a$ s.t. $a = \text{assign}(c, id, T, \vec{i}, \vec{p})$. We denote it with $T \triangleleft X$.

The two latter definitions allow to state a new further definition of $\Phi(s)$ and $\Psi(s)$, whose range can be limited to a specific task T .

Definition 6. Given a specific task T , a T -limited physical reality $\Phi|_T(s)$ is the set of that data fluents $X_{\varphi,y}$ (where y ranges over $\{1..m\}$) such that $T \triangleright X_{\varphi,y}$ or $T \triangleleft X_{\varphi,y}$. We denote these fluents as $X_{\varphi|T}$. Hence, $\Phi|_T(s) = \bigcup_{y=1..m} \{X_{\varphi,y|T}\}$ and $\Phi|_T(s) \subseteq \Phi(s)$.

Definition 7. Given a specific task T , a T -limited expected reality $\Psi|_T(s)$ is the set of that expected fluents $X_{\psi,y}$ (where y ranges over $\{1..m\}$) such that $T \triangleright X_{\psi,y}$ or $T \triangleleft X_{\psi,y}$. We denote these fluents as $X_{\psi|T}$. Hence, $\Psi|_T(s) = \bigcup_{y=1..m} \{X_{\psi,y|T}\}$ and $\Psi|_T(s) \subseteq \Psi(s)$.

From definitions 6 and 7, the following ones stem :

Definition 8. Let T_1, \dots, T_n all tasks defined in the SitCalc theory. A physical reality $\Phi(s)$ is the union of all T -limited physical realities that hold in situation s : $\Phi(s) = \bigcup_{i=1..n} \Phi|_{T_i}(s)$.

Definition 9. Let T_1, \dots, T_n all tasks defined in the SitCalc theory. An expected reality $\Psi(s)$ is the union of all T -limited expected realities that hold in situation s : $\Psi(s) = \bigcup_{i=1..n} \Psi|_{T_i}(s)$.

Now, the predicate *Relevant* can be easily refined in a way that focuses on a specific task T :

$$\text{Relevant}_T(\delta, s) \equiv \neg \text{SameState}(\Phi|_T(s), \Psi|_T(s)) \quad (3.24)$$

Example 3.5.2. Let us consider again the example described in our case study. The task $go(loc00, loc33)$ was assigned to actor $act1$, which has reached the wrong location $loc03$ (cf. Fig. 1.3(b)). Therefore, $act1$ is now located in a different position than the desired one, and s/he is out of the optimal network range. After the *release* action has been executed, the fluent At_{φ} takes the value $loc03$, while the expected output $loc33$ is stored in the fluent At_{ψ} . This generates a discrepancy between $\Phi|_{go}(s)$ and $\Psi|_{go}(s)$. This means that $\text{Relevant}_{go}(\delta, s)$ holds, and the main process δ needs to be adapted.

In Fig. 3.8 we show how we have concretely coded the continuous planning approach with IndiGolog. Like in the plan-based adaptation approach, before starting the execution of the process δ , the PMS builds the PDDL representation of each task defined in the SitCalc theory and sends it to a whatever external planner that implements the POP algorithm.

The main procedure involves four concurrent programs in priority. The interrupt at lowest priority, that consists just in waiting, is activated only if the process is still not finished, but for some reason, it can not progress in its execution (e.g., the PMS is unable to find any service providing all the capabilities required by a specific task

```

PROC Main()
1   $\langle RealityChanged \wedge \neg Finished \rightarrow [Monitor()] \rangle$ .
2   $\langle Recovered \wedge \neg Finished \rightarrow [UpdateProcess()] \rangle$ .
3   $\langle \neg Finished \rightarrow [Process(); finish] \rangle$ .
4   $\langle \neg Finished \rightarrow [wait] \rangle$ ;

PROC UpdateProcess()
1  receive(Planner,  $plan_h$ ).
2  Convert( $plan_h, \delta_h$ ).
3  Update( $\delta_0, (\delta_0 || \delta_h)$ ).
4  resetReality.
5  resetRecovery.

PROC Monitor()
1  IndiPOP().
2  resetRealityChanged.

PROC IndiPOP()
1   $Start = \Phi(s) \cup \{\bigcup_{i=1..k, j=1..n} \{Capable(c_i, t_j)\}\} \cup \{\bigcup_{i=1..k} \{Free(c_i, s)\}\}$ .
2   $Finish = \Psi(s)$ .
3  receive(Planner,  $plan_p$ ).
4  send(RemoveConflicts, [ $plan_p, Start, Finish$ ]).
5  receive(RemoveConflicts,  $plan_u$ ).
6  send(Planner, [ $plan_u, Start, Finish$ ]).

```

Figure 3.8. The IndiGolog main procedure customized for the continuous planning approach.

to be executed). If the fourth interrupt is activated, the control passes back to the process designer, which can manually manage the situation.

Then, the system runs the actual IndiGolog program representing the process to be executed (the procedure **Process**). This procedure relies, in turn, on procedure **ManageExecution**, which includes task assignment, start signaling, acknowledgment of completion, and final release (cf. Section 3.3.1). The monitor, which runs at higher priority, is in charge of monitoring changes in the environment and adapting accordingly. The first step in procedure **Monitor** checks whether fluent *RealityChanged* holds true, meaning that a service has terminated the execution of a task or an exogenous (unexpected) action has occurred in the system. Basically, the procedure **Monitor** is enabled when the physical or the expected reality (or both) change.

If it happens, the monitor calls the procedure **IndiPOP**, whose purpose is to manage the execution of the external planner by updating its initial states and expected goals according with changes in the two realities. **IndiPOP** first builds the two sets *Start* (the initial state) and *Finish* (the goal), by making them equal respectively to $\Phi(s)$ and $\Psi(s)$. As far as concerns the initial state, it will include, for each task t and service c defined in SitCalc theory, the values of *Capable*(c, t) and of *Free*(c, s) in addition to the values of data fluents.

Then **IndiPOP** catches the partial plan $plan_p$ (that has the form of a set of partial ordering constraints between tasks; it is empty if no failure has happened yet) built till that moment by the external planner and updates it with the new sets

Start and Finish. Such updating finds something about $plan_p$ that needs fixing in according with the new realities. Since $plan_p$ has been built working on old values of the two realities, it is possible that some ordering constraints between tasks are not valid anymore. This causes the generation of some conflicts, that need to be deleted by $plan_p$ through the external procedure **RemoveConflicts**. Basically, **IndiPOP** can be seen as a conflict-removal procedure that revises the partial recovery plan to the new realities. At this point, $plan_u$ (that is, $plan_p$ just updated, i.e., without conflicts) is sent back to the external planner together with the sets Start and Finish. The external planner can now restore its planning procedure.

Note that if the predicate *Relevant*(s) holds, meaning that a misalignment between the two realities exists, the PMS tries to continue with its execution. In particular, every T_i whose T-limited expected reality $\Psi|_{T_i}(s)$ is different from the T-limited physical reality $\Phi|_{T_i}(s)$ could not anymore proceed with its execution. However, every task T_j not affected by the deviation can advance without any obstacle. Once sent the sets of fluents composing the two realities to the external planner, the monitor resets the fluent *RealityChanged* to *false*, and the control passes to the process of interest (i.e., program **Process**), that may again execute/advance.

When the external planner finds a recovery plan that can align physical and expected reality, the fluent *Recovered* is switched to *true* and the procedure **UpdateProcess** is enabled. Now, after receiving the recovery process δ_a from the planner, the PMS updates the original process δ to a new process δ'' that, respect to its predecessor, has a new branch to be executed in parallel; such branch is exactly δ_a . It contains all that tasks able to repair the physical reality from the discrepancies (i.e. to unblock all that tasks stopped in δ because their preconditions did not hold). Note that when δ_a is merged with the original process δ , the two realities are still different from each others. Therefore, the PMS makes them equal by forcing $\Psi(s)$ to the current value of $\Phi(s)$. This because the purpose of δ_a , after that all recovery actions have been executed, is to turn the current $\Phi(s)$ into $\Psi(s')$, where s' is that situation reached after the execution of recovery actions.

Let us now formalize the concept of *strongly consistency* for a process δ .

Definition 10. *Let δ a process composed by n tasks T_1, \dots, T_n . δ is **strongly consistent** iff:*

- *Given a specific task T and an input I , $\nexists c, c', \vec{p}, \vec{p}', \vec{q}, \vec{q}', a, a'$ s.t.
 $a = \text{release}(c, T, I, \vec{p}, \vec{q}) \wedge a' = \text{release}(c', T, I, \vec{p}', \vec{q}') \wedge (\vec{p} \neq \vec{p}')$.*
- *$\forall y \in 1..m, \nexists (T_i, T_k)_{i \neq k}$ s.t. $(T_i \triangleright X_{\varphi, y} \wedge T_k \triangleright X_{\varphi, y})$.*

Intuitively, a process δ is strongly consistent if a specific task, executed on a given input, cannot return different values for its expected output; moreover, the above condition holds if do not exist two different tasks that affect the same fluent. For strongly consistent processes, we can state the concept of *goal* :

Definition 11. *Given a strongly consistent process δ , composed by n tasks T_1, \dots, T_n , the goal of δ can be defined as the set of all expected fluents $X_{\psi, y}$ that are affected by T_1, T_2, \dots, T_n . Hence, $Goal(\delta) = \{X_{\psi, y} \text{ s.t. } \exists i_{1..n}. (T_i \triangleright X_{\psi, y})\}$.*

After a recovery procedure δ_a , $Goal(\delta) \subseteq Goal(\delta || \delta_a)$, since the recovery procedure can introduce new tasks with respect to the original process δ . Anyway, the original $Goal(\delta)$ is preserved also after the adaptation procedure.

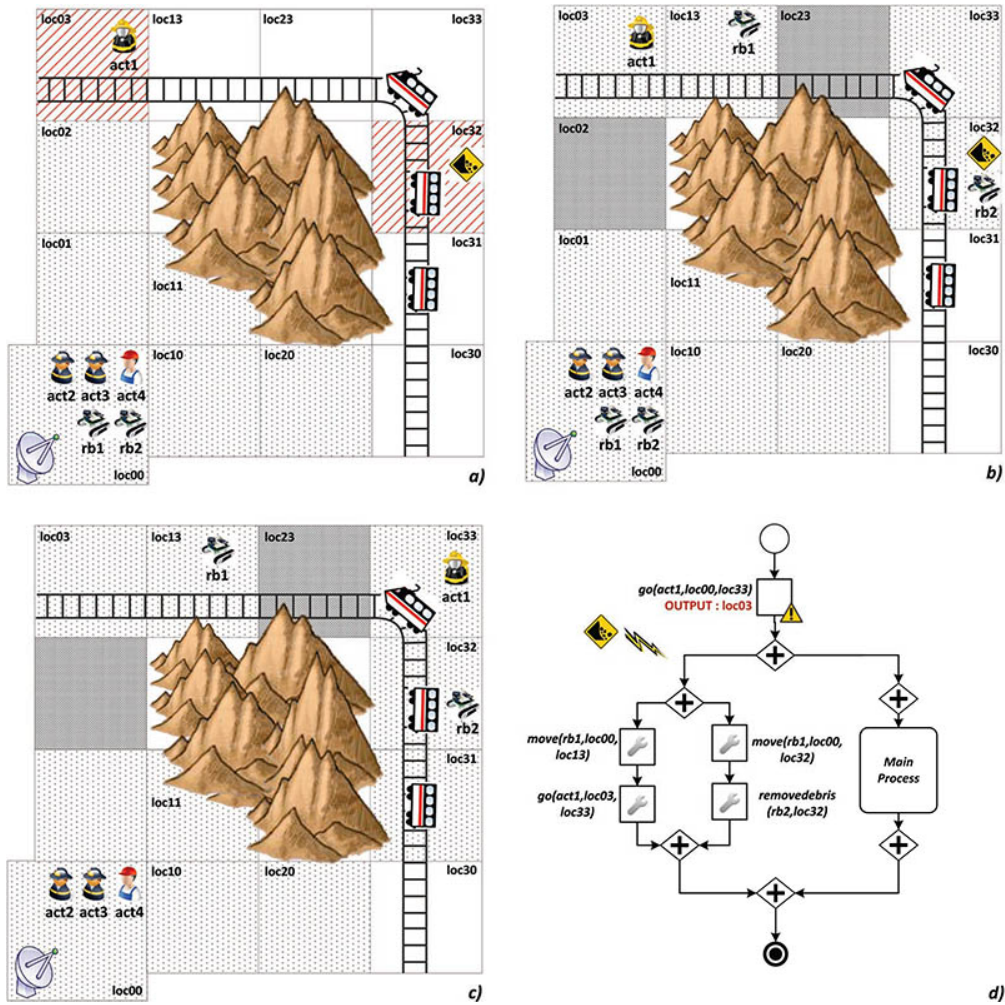


Figure 3.9. An example of adaptation through the continuous planning approach.

Theorem 4 (Termination). *Let δ be a strongly consistent process composed by a finite number of tasks T_1, \dots, T_n . If δ does not contain while and iteration constructs, and the number of exogenous events is finite, then the core procedure of IndiGolog PMS terminates.*

Proof. (Sketch) We have to consider every case in which one of the two reality can change, since only in such a case the PMS updates the recovery plan under construction and restarts the external planner with new values of initial state and goals.

- *Case 0 : No Failures.* We know that both $\Phi(s)$ and $\Psi(s)$ can change when a task T ends its execution. As the number of tasks composing the original process δ is finite, if no internal failures or exogenous events occur (meaning that each task terminates its execution by returning the expected output), the process terminates without any need to compute a recovery procedure.
- *Case 1 : Internal Failures.* Again, since the number of tasks composing the original process δ is finite, also the number of *internal failures* that could

happen will be finite. In the worst case, in which $\delta = T_1 || T_2 || \dots || T_n$, if every task fails during its execution, we will have exactly n internal failures, meaning that $\Phi(s)$ and $\Psi(s)$ change n times. In such a case, no task can more proceed in the execution, and the external planner can terminate (with a recovery plan δ_a or with a failure) without any more alteration of its initial situation and goals.

- *Case 2 : Exogenous Events.* As the number of exogenous events that affect the execution of δ is finite, when the very last exogenous event has occurred, the only way to change realities is either by the standard execution of δ , or by internal failures, falling respectively into *Case 0* or *Case 1*.

□

We want to underline that the termination cannot be guaranteed if δ contains loops or iteration, since potentially the two realities could indefinitely change. The same is true if the number of exogenous events is unbounded. In such a case, there exists the possibility that **IndiPOP** updates continuously the external planner, that would not be able to terminate by returning the expected recovery plan.

Example 3.5.3. Let us suppose that during the synthesis of the recovery plan needed for dealing with the exception arose in the previous example 3.5.2, an exogenous event *rockSlide(loc32)* is captured by the PMS. It aims at alerting about a rock slide collapsed in location *loc32* (cf. Fig. 3.9(a)). Hence, we have different values for $Status_\varphi(loc32, s') = debris$ and $Status_\psi(loc32, s') = ok$. Again, the PMS invokes the external planner by obtaining the partial plan $plan_p$ built till that moment and verifies if it needs to be fixed according with new values of the two realities. If no conflicts are individuated, the PMS sends back $plan_p$ to the planner together with the information about the initial state and the goal, updated to situation s' . When the planner ends its computation, it returns a recovery process δ_a that includes two concurrent branches. The left branch of δ_a instructs robot *rb1* to reach a position where actor *act1* is again connected to the network, and can finally reach its original expected destination *loc33* (cf. Fig. 3.9(c)). In parallel, robot *rb2* can reach location *loc32* and remove debris. Note that δ_a can be executed in concurrency with δ (see the right-hand side of Figure 3.9(d)) by preserving its original goal.

3.6 Conclusion

In this chapter we have presented the formal foundations of our general approach based on formalizing processes with Situation Calculus and **IndiGolog**, on detecting exceptions with execution monitoring and on adapting automatically a dynamic process through classical planning techniques. Such an approach is (i) practical, since it relies on well-established planning techniques, and (ii) does not require the definition of the adaptation strategy in the process itself (as most of the current approaches do).

We also gave details on three different techniques for adapting a faulty process instance. The built-in adaptation approach, based on the Σ search construct

provided by IndiGolog, allows to incorporate adaptation features directly into the PMS. However, the computation of a recovery plan with this approach is very inefficient, since it is based on a breadth first search algorithm that results in an exponential explosion in the size of the inputs/services.

The plan-based technique for adapting processes is instead very efficient, since it delegates the building of the plan to state-of-the-art external planners. Specifically, we used the LPG-td planner [45] for the synthesis of the recovery plans, and in Chapter 5 we clearly demonstrate the feasibility of this choice by showing some interesting experimental tests. The drawback of this approach lies in the intrinsic assumptions of classical planning (determinism in the effects, model completeness, etc.), that could be too restrictive to address complex problems.

The continuous planning approach seems very promising for reducing the overall response time during adaptation. The strength of the approach lies in the ability to incorporate execution feedback directly into the plan, without blocking the execution of the main process. In fact, during the plan synthesis, if a positive event occurs (such as an external event or a task whose effects are to “adjust” the compromised situation), the system is able to take advantage of such an opportunity without the need of synthesizing a new plan. However, even though our intent is to make the planning process very responsive, there still remains a synchronization process between planning and execution, which can require a significant response time. Another issue concerns the augmented response time needed by the planner when more than one exception is treated at the same time. In such a case, the recovery plan could be constituted by n recovery sub-processes independent one from another, and the overall time depends by the number of steps required for finding the longest one (in terms of recovery activities).

Chapter 4

The SmartPM System

The SmartPM approach described in Chapter 3 has shown that a PMS supporting process adaptation at run-time requires an integrated approach that covers the modeling, execution and monitoring stages of process life-cycle. The combination of procedural and imperative models with declarative elements, along with the exploitation of techniques from the field of artificial intelligence (AI) such as planning algorithms and tools, is required for increasing the ability of a PMS of supporting dynamic processes.

To this end, in this chapter we aim at presenting the overall architecture of SmartPM and at giving technical details about its atomic components. In Section 4.1, we show how software components of SmartPM have been organized into multiple logical layers, which correspond to the system's main features. Then, in Section 4.2 we present the IndiGolog platform we used for building our PMS and the execution monitor. In Section 4.3 we describe SmartML, which combines a modeling formalism for representing the information of the contextual scenario linked to a specific dynamic process, and a graphical tool (specifically, Eclipse BPMN¹) for designing the control flow of the process. We also show how a dynamic process formalized through SmartML is automatically translatable in situation calculus and IndiGolog readable formats and is therefore ready for being executed by SmartPM. Finally, in Section 4.4 we give some high level detail on the translation algorithms used for building planning domains and problems starting from SmartML/IndiGolog specifications, and in Section 4.5, we show some screenshots of the SmartPM system in action.

4.1 System Architecture

Our approach to integration of process execution and planning for providing automatic adaptation features at run-time relies on three main architectural layers as shown in Fig. 4.1.

The **Presentation Layer** has a twofold purpose. On one hand, it allows a human process designer to define a dynamic process and all the contextual information concerning the scenario in which the process will be enacted. To this end, we provide a clarifying descriptive tool called *SmartPM Definition Tool*. It is a GUI-based tool that can be used by a process designer to build a process specification

¹<http://www.eclipse.org/modeling/mdt/?project=bpmn2>

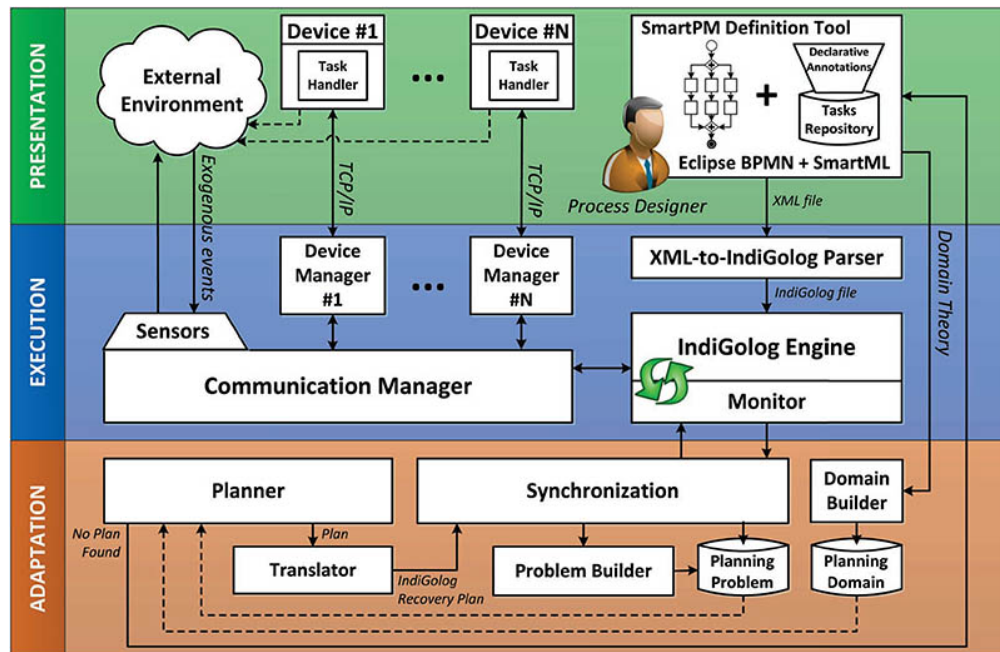


Figure 4.1. The architecture of SmartPM.

defined according to the SmartPM Modeling Language (a.k.a. SmartML) presented in Section 4.3.1. The language allows to specify a so-called *Domain Theory*, i.e., to clearly represent the contextual properties reflecting a dynamic scenario. Moreover, through SmartML a process designer can build a repository of tasks defined in a declarative way, by explicitly providing tasks pre-conditions and effects based on the domain theory previously defined. Finally, the process designer can build graphically the control flow of a dynamic process through *Eclipse BPMN*², a graphical editor that allows to specify business processes using the BPMN 2.0 notation [7]. Note that the set of tasks composing the control flow of the dynamic process must be selected from the repository of tasks specified through SmartML. In Section 4.3 we will show how to build a dynamic process through the SmartPM Definition Tool, that integrates the graphical features of Eclipse BPMN with the SmartML declarative language. The outcome of the process design activity will be a complete XML-encoded process specification to be passed to the IndiGolog engine.

On the other side, the Presentation Layer includes every real world device that may interact with the IndiGolog engine. In order to manage such an interaction, there is the need to installing and configuring a *Task Handler* module on top of every device. The Task Handler is an interactive GUI-based software application that - during process execution - supports the visualization of assigned tasks and allows each participant to start task execution and notify task completion by selecting an appropriate outcome. A task is seen by the Task Handler as a black-box activity, and the only relevant information captured are the starting and the completion of the task itself (together with the respective outcomes). Finally, since our SmartPM system is mainly a proof-of-concept implementation of an adaptive PMS, we simulate the

²<http://eclipse.org/bpmn2-modeler/>

External Environment as a software module that sends asynchronously exogenous events that may prevent the correct execution of the dynamic process.

The **Execution Layer** is in charge of managing and coordinating the execution of dynamic processes. It performs task assignments and it manages and stores information about services and tasks involved in process executions, tasks to be completed and variables/data modified during task executions. The Execution Layer is basically a customized version of the IndiGolog platform, whose details are provided in Section 4.2.

A dynamic process built through Eclipse BPMN and annotated with SmartML is taken as input from the *XML-to-IndiGolog Parser* component, which translates this specification in Situation Calculus and IndiGolog readable formats, in order to make the process executable by the IndiGolog engine.

The *IndiGolog Engine* provides a proper execution engine that manages the process routing and decides which tasks are enabled for execution, by taking into account the control flow, the value of predicates and preconditions and effects of each task. Before a process starts its execution, the IndiGolog engine builds its physical reality Φ_s by taking the initial context from the environment, and the expected reality Ψ_s , which initially is equal to Φ_s .

Once a task is ready for being assigned, a component named *Communication Manager* is in charge of assigning it to a proper service (which may be a human actor, a robot, a software application, etc.) that is available (i.e., free from any other task assignment) and that provides all the required capabilities for task execution. Every step of the task life cycle - ranging from the assignment to the release of a task - requires an interaction between the Communication Manager and the devices. For each real world device, the Communication Manager generates a separate *Device Manager*, which is a software component that is able to interact with the Task Handlers deployed on the devices. Each device manager establishes a communications channel with the associated device by using TCP/IP stream sockets. Such an interaction is mainly intended for notifying the corresponding device of actions performed by the IndiGolog engine as well as for notifying the engine of actions executed by the Task Handlers of the corresponding device. Finally, the Communication Manager allows also to catch exogenous events coming from the environment and, when the process terminates, to notify the process completion to the devices (again, through the Device Managers).

The *Monitor* component, which interacts continually with the IndiGolog engine, is in charge of monitoring contextual data in order to identify changes, modifications or events which may affect process execution, and notify them to the Adaptation layer. Specifically, at each execution step - i.e., when the ending of a task or an exogenous event has turned Φ_s into Φ_{s+1} - the monitor checks if the new situation $s + 1$ can be classified as relevant (cf. Equation 3.20). If this is the case, the monitor collects the physical reality Φ_{s+1} , the expected reality Ψ_{s+1} and sends them to the synchronization component. In a nutshell, the Monitor component decides whether adaptation is needed.

The **Adaptation Layer** is in charge of reacting to undesired or unforeseen events which may invalidate process execution. The *Synchronization* component acts as unique entry point for incoming notifications from the Execution Layer, in order to adapt process execution through specific recovery/adaptation techniques. It enforces

synchronization between the IndiGolog engine, the monitor and the planner. Every time it receives from the monitor the two realities, it builds a corresponding planning problem in PDDL (through the *Problem Builder* component), by converting the physical reality Φ_s into the initial state and the expected reality Ψ_s into the goal.

The *Planner* component is invoked when the Synchronization component builds a new planning problem. In addition to the initial state and the goal, the Planner needs a specification for the planning domain too (that is, a PDDL specification with tasks and predicates). For this purpose, the *Domain Builder* component translates the domain theory defined in SmartML in a PDDL planning domain readable by the Planner. In the SmartPM system, we synthesize our recovery plans through the LPG-td planner [45] (Local search for Planning Graphs). It is a planner based on local search and planning graphs that handles PDDL 2.2 domains and can produce good quality plans in terms of one or more criteria. In order to invoke the planner, the Synchronization component needs to specify the value of three parameters indicating: (i) a file containing a set of PDDL operators (i.e., the planning domain, provided by the Domain Builder), (ii) a file containing a planning problem (i.e., the initial state and goal of the problem, provided by the Problem Builder) formalized in PDDL and (iii) a running mode, which is either “speed” (for devising sub-optimal solutions) and “quality” (for devising optimal solutions). More details about the planner are given in Section 5.1.

Finally, when a plan satisfying the goal is found, it is sent back to a *Translator* component, that converts it in a readable format for the IndiGolog engine and passes it back to the Synchronization component. The Synchronization component combines the faulty process instance δ' with the recovery plan δ_a just built, and obtains the adapted process $\delta'' = (\delta_a; \delta')$ to be executed by the IndiGolog engine. If the Planner is not able to find any recovery plan for a specific deviation, the control is given back to the process manager.

4.2 The IndiGolog Platform

The IndiGolog platform is a logic-programming implementation of IndiGolog that allows the incremental execution of high-level Golog-like programs. Part of this section is a summary of the work published in [25], that we customized to our needs in collaboration with the creators of IndiGolog. Although most of the code is written in vanilla Prolog, the overall architecture is written in the well-known open source SWI-Prolog³ [144]. SWI-Prolog provides flexible mechanisms for interfacing with other programming languages such as Java or C, allows the development of multi-threaded applications, and provides support for socket communication and constraint solving.

Generally speaking, the IndiGolog implementation provides an incremental interpreter of high-level programs as well as a framework for dealing with the real execution of these programs on concrete platforms or devices. This amounts to handling the real execution of actions on concrete devices (e.g., a PDA), the collection of sensing outcome information (e.g., retrieving some sensor’s output), and the detection of exogenous events happening in the world. To that end, the architecture

³Available at <http://www.swi-prolog.org/>

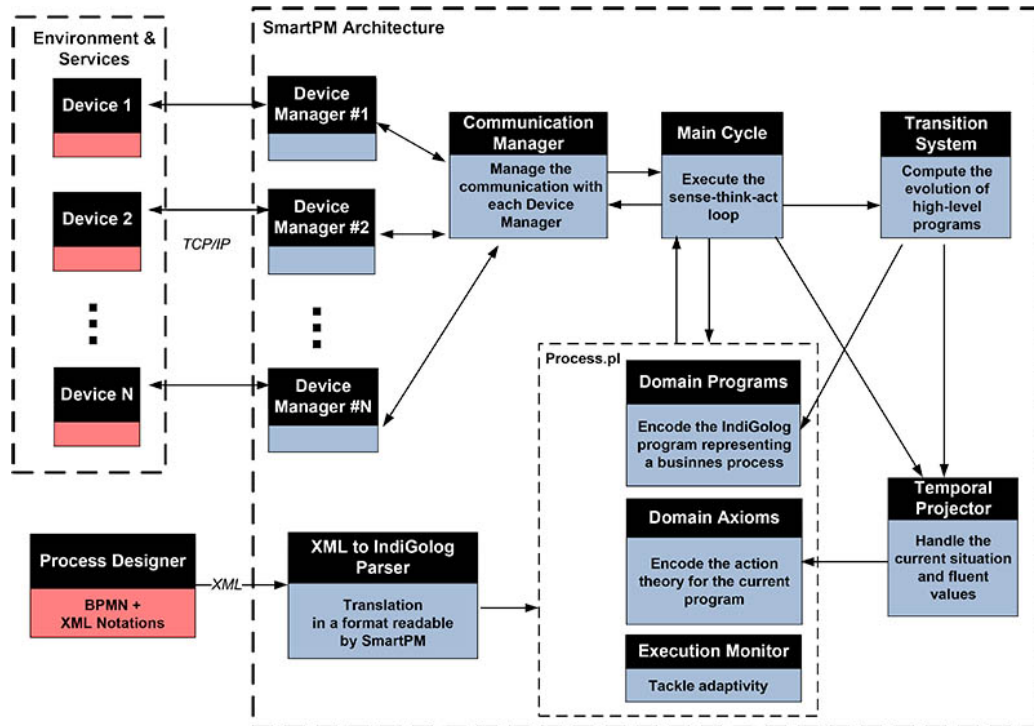


Figure 4.2. The architecture of the IndiGolog platform.

is modularly divided into six parts, namely, (i) the top-level main cycle; (ii) the language semantics; (iii) the temporal projector; (iv) the communication manager; (v) the set of device managers; and finally (vi) the domain application. The first four modules are completely domain independent, whereas the last two are designed for specific domain(s). The architecture is depicted in Fig. 4.2.

4.2.1 The Top-level Main Cycle and Language Semantics

The IndiGolog platform codes the *sense-think-act* loop well-known in the agent community [65]:

1. check for exogenous events that have occurred;
2. calculate the next program step; and
3. if the step involves an action, *execute* the action.

While executing actions, the platform keeps updated an history, which is the sequence of actions performed so far.

The main predicate of the main cycle is `indigo/2`; a goal of the form `indigo(E,H)` states that the high-level program `E` is to be executed online at history `H`.

The first thing the main cycle does is to assimilate all exogenous events that have occurred since the last execution step. After all exogenous actions have been assimilated and the history progressed as needed, the main cycle goes on to actual executing the high-level program `E`. First, if the current program to be executed is terminating in the current history, then the top-level goal `indigo/2` succeeds.

Otherwise, the interpreter checks whether the program can evolve a single step by relying on predicate `trans/4` (explained below). If the program evolves without executing any action, then the history remains unchanged and we continue to execute the remaining program from the same history. If, however, the step involves performing an action, then this action is executed and incorporated into the current history, together with its sensing result (if any), before continuing the execution of the remaining program.

As mentioned above, the top-level loop relies on two central predicates, namely, `final/2` and `trans/4`. These predicates implement relations *Trans* and *Final*, giving the single step semantics for each of the constructs in the language. It is convenient, however, to use an implementation of these predicates defined over histories instead of situations. Indeed, the constructs of the IndiGolog interpreter never treat about situations but they are always assuming to work on the current situation. So, for example, these are the corresponding clauses for sequence (represented as a list), tests, nondeterministic choice of programs, and primitive actions:

```
final([E|L],H) :- final(E,H), final(L,H).
trans([E|L],H,E1,H1) :- final(E,H), trans(L,H,E1,H1).
trans([E|L],H,[E1|L],H1) :- trans(E,H,E1,H1).

final(ndet(E1,E2),H) :- final(E1,H) ; final(E2,H).
trans(ndet(E1,E2),H,E,H1) :- trans(E1,H,E,H1).
trans(ndet(E1,E2),H,E,H1) :- trans(E2,H,E,H1).

trans(?P),H,[],H) :- eval(P,H,true).
trans(E,H,[],[E|H]) :- action(E), poss(E,P), eval(P,H,true).
/* Obs: no final/2 clauses for action and test programs */
```

These Prolog clauses are almost directly “lifted” from the corresponding axioms for *Trans* and *Final*. Predicates `action/1` and `poss/2` specify the actions of the domain and their corresponding precondition axioms; both are defined in the domain axiomatization (see below). More importantly, `eval/3` is used to check the truth of a condition at a certain history, and is provided by the temporal projector, described next.

The naive implementation of the search operator would deliberate from scratch at every point of its incremental execution. It is clear, however, that one could do better than that, and cache the successful plan obtained and avoid planning in most cases:

```
final(search(E),H) :- final(E,H).
trans(search(E),H,path(E1,L),H1) :-
    trans(E,H,E1,H1), findpath(E1,H1,L).

/* findpath(E,H,L): solve (E,H) and store the path in list L */
/* L = list of configurations (Ei,Hi) expected along the path */
findpath(E,H,[E,H]) :- final(E,H).
findpath(E,H,[E,H|L]) :- trans(E,H,E1,H1), findpath(E1,H1,L).
```

So, when a search block is solved, the whole solution path found is stored as the sequence of configurations that are expected. If the actual configurations match, then steps are performed without any reasoning (first `final/2` and `trans/4` clauses for program `path(E,L)`). On the other hand, if the actual configuration does not match the one expected next, for example, because an exogenous action occurred and the history thus changed, re-planning is performed to look for an alternative path (code not shown).

4.2.2 The Temporal Projector

The temporal projector is in charge of maintaining the agent's beliefs about the world and evaluating a formula relative to a history. The projector module provides an implementation of predicate `eval/3`: goal `eval(F,H,B)` states that formula `F` has truth value `B`, usually `true` or `false`, at history `H`.

Predicate `eval/3` is used to define `trans/4` and `final/2`, as the legal evolutions of high-level programs may often depend on what things are believed true or false.

We assume then that users provide definitions for each of the following predicates for fluent `f`, action `a`, sensing result `r`, formula `w`, and arbitrary value `v`:

```

fun_fluent(f) f is a functional fluent;
rel_fluent(f) f is a functional fluent;
prim_action(a) a is a ground action;
init(f,v) v is the value for fluent f in the starting situation;
poss(a,w) it is possible to execute action a provided formula w is known to be true;
causes_val(a,f,v,w) action a affects the value of f

```

Formulas are represented in Prolog using the obvious names for the logical operators and with all situations suppressed; histories are represented by lists of the form $o(a,r)$ where a represents an action and r a sensing result. We will not go over how formulas are recursively evaluated, but just note that there exists a predicate (*i*) `kTrue(w,h)` is the main and top-level predicate and it tests if the formula w is at history h . Finally, the interface of the module is defined as follows:

```

eval(F,H,true) :- kTrue(F,H).
eval(F,H,false) :- kTrue(neg(F),H).

```

4.2.3 The Communication Manager

The IndiGolog system was meant for being used with concrete agent/robotic platforms, as well as with software/simulation environments. To this end, the online execution of IndiGolog programs must be linked with the external world. To this end, the *communication manager* (CM) provides a complete interface with all the external devices, platforms, and real-world environments that the application needs to interact with. In turn, each external device or platform that is expected to interact with the application (e.g., a robot, a software module, or even a user interface) is assumed to

have a corresponding *device manager*, a piece of software that is able to talk to the actual device, instruct it to execute actions, as well as gather information and events from it. The device manager understands the “hardware” of the corresponding device and provides a high-level interface to the CM. It provides an interface for the execution of exogenous events (e.g., *assign*, *start*, etc.), the retrieval of sensing outcomes for actions, and the occurrence of exogenous events (e.g., *photoLost* as well as *finishedTask* and *readyToStart*). Because actual devices are independent of the IndiGolog application and may be in remote locations, device managers are meant to run in different processes and, possibly, in different machines; they communicate then with the CM via TCP/IP sockets. The CM, in contrast, is part of the IndiGolog agent architecture and is tightly coupled with the main cycle. Still, since the CM needs to be open to the external world regardless of any computation happening in the main cycle, the CM and the main cycle run in different (but interacting) threads, though in the same process and Prolog run-time engine.⁴ So, in a nutshell, the CM is responsible of executing actions in the real world and gathering information from it in the form of sensing outcome and exogenous events by communicating with the different device managers.

More concretely, given a domain high-level action (e.g., *assign(Srv, Id, Task, Inputs, ExOutputs)*), the CM is in charge of: (i) deciding which actual “device” should execute the action; (ii) ordering its execution by the device via its corresponding device manager; and finally (iii) collecting the corresponding sensing outcome. To realize the execution of actions, the CM provides an implementation of *exec/2* to the top-level main cycle: *exec(A, S)* orders the execution of action *A*, returning *S* as its sensing outcome.

When the system starts, the CM starts up all device managers required by the application and sets up communications channels to them using TCP/IP stream sockets. Recall that each real world device or environment has to have a corresponding device manager that understands it.

After this initialization process, the CM enters into a *passive mode* in which it asynchronously listens for messages arriving from the various devices managers. This passive mode should allow the top-level main cycle to execute without interruption until a message arrives from some device manager.

In general, a message can be an exogenous event, a sensing outcome of some recently executed action, or a system message (e.g., a device being closed unexpectedly). The incoming message should be read and handled in an appropriate way, and, in some cases, the top-level main cycle should be notified of the occurred event.

4.2.4 The Domain Application

From the user perspective, probably the most relevant aspect of the architecture is the specification of the domain application. Any domain application must provide:

1. An *axiomatization of the dynamics of the world*. Such axiomatization would depend on the temporal projector to be used.

⁴SWI-Prolog provides a clean and efficient way of programming multi-threaded Prolog applications.

2. One or more *high-level agent programs* that will dictate the different agent behaviors available. In general, these will be IndiGolog programs.
3. All the necessary *execution information* to run the application in the external world. This amounts to specifying which external devices the application relies on (e.g., the device manager for the ER1 robot), and how high-level actions are actually executed on these devices (that is, by which device each high-level action is to be executed). Information on how to translate high-level symbolic actions and sensing results into the device managers' low-level representations, and vice-versa, could also be provided.

4.3 The SmartPM Definition Tool

A process modeling language provides appropriate syntax and semantics to precisely specify business process requirements, in order to support automated process verification, validation, simulation and automation. One of the main obstacles in applying AI techniques to real problems is the difficulty to model the domains. Usually, this requires that people that have developed the AI system carry out the modeling phase since the representation depends very much on a deep knowledge of the internal working of the AI tools. On the contrary, during the realization of the SmartPM system we took care of the above aspect, and we worked on the definition of a language (named SmartML) that allow non-experts entering knowledge on processes through a user-friendly interface. Such knowledge, together with the business process to be enacted, is automatically translated in Situation Calculus, IndiGolog and PDDL.

In this section, we aim at describing the main components of the SmartPM Definition Tool. Specifically, we first present SmartML, our declarative language used for representing tasks and contextual data linked to a pervasive scenario, and *Eclipse BPMN*, a graphical editor that allows to specify business processes using the BPMN 2.0 notation [7]. Then, we focus on the XML-to-IndiGolog Parser component, that translates a SmartML specification in a SitCalc Theory (which defines the initial situation and the set of data fluents and available actions with their pre- and post-conditions) and the BPMN control flow of the process in a IndiGolog program corresponding to the process to be executed.

4.3.1 The SmartML Modeling Language

The synthesis of a dynamic process requires a tight integration of process activities and contextual data in which the process is embedded in. The context is represented in the form of a *Domain Theory* D , that involves capturing a set of tasks $t_i \in T$ (with $i \in 1..n$) and supporting information, such as the people/agents that may be involved in performing the process (roles or participants), the data and so forth.

Tasks are collected in a specific repository, and each task can be considered as a single step that consumes input data and produces output data. Data are represented through some ground atomic terms $v_1[y_1], v_2[y_2], \dots, v_m[y_m] \in V$ that range over a set of tuples (i.e., unordered sets of zero or more attributes) y_1, y_2, \dots, y_m of *data objects*, defined over some *data types*. In short, a data object depicts

an entity of interest. Some data types are pre-specified and used for representing the *resource perspective* of the process.

For example, in our scenario we need to define data objects for representing participants (e.g., data type $Participant = \{act1, act2, act3, act4, rb1, rb2\}$) and capabilities (e.g., data type $Capability = \{extinguisher, movement, \dots hatchet\}$).

Resource Perspective :

$Participant = \{act1, act2, act3, act4, rb1, rb2\}$
 $Capability = \{movement, hatchet, camera, gprs, extinguisher, battery, digger, powerpack\}$

The data types $Participant$ and $Capability$ are already pre-defined for being used in the framework, and the Process Designer is only required to provide values (i.e., to associate data objects) to the above types.

On the contrary, other data types need to be defined for describing the contextual scenario in which the process will be embedded. In our example, we may need of a data type $Location_type = \{loc00, loc10, \dots loc33\}$ for representing locations in the area. Moreover, a data type $Status_type = \{ok, fire, debris\}$ may be defined for describing if a specific location is on fire or buried by debris.

User-Defined Data Types :

$Location_type = \{loc00, loc10, loc20, loc30, loc01, loc11, loc02, loc03, loc13, loc23, loc31, loc32, loc33\}$
 $Status_type = \{ok, fire, debris\}$

Under this representation, we consider possible values of a data type as constant symbols that univocally identify data objects in the scenario of interest. Each tuple y_j may contain one or more data objects belonging to different data types. The domain $dom(v_j[y_j])$ over which a term is interpreted can be of various types:

- $Boolean_type$: $dom(v_j[y_j]) = \{true, false\}$;
- $Integer_type$: $dom(v_j[y_j]) = \{x \dots y\}$ s.t. $x, y \in \mathbb{Z} \wedge (x \leq y)$;
- $Functional$: the domain contains a fixed number of data objects of a designated type.

The data type $Boolean_type$ (and its respective “objects” $true$ and $false$) and the type $Integer_type$ are already pre-specified. However, since integer numbers form a countably infinite set, there is the need to set a lower and an upper bound to specify which finite subset of integers is relevant for the case to deal with. In the below example, we are considering the subset of the integers from 0 to 30.

Pre-Defined Data Types :

$Boolean_type = \{true, false\}$
 $Integer_type = \{0, 1, 2, 3, 4, \dots, 30\}$

Terms can be used to express properties of domain objects (and relations over objects), and argument types of a term - taken from the set of data types previously

defined⁵ - represent the finite domains over which the term is interpreted. Again, some terms are pre-specified for being used in the framework. For example, since each task has to be assigned to a participant that provides all of the skills required for executing that task, there is the need to consider the participants “capabilities”. This can be done through a boolean term $provides[prt : Participant, cap : Capability]$ that is *true* if the capability *cap* is provided by *prt* and *false* otherwise. At the same way, a boolean term $requires[task : Task, cap : Capability]$ is needed for specifying which capabilities are required for executing a specific task⁶.

Pre-Defined Terms :

```
provides[prt:Participant,cap:Capability] = (bool:Boolean_type)
requires[task:Task,cap:Capability] = (bool:Boolean_type)
```

Moreover, we may need boolean terms for indicating if people have been evacuated from a location (e.g., $evacuated[loc : Location_type] = (bool : Boolean_type)$), integer terms for representing the battery charge level of each robot (e.g., $batteryLevel[prt : Participant] = (int : Integer_type)$) or for indicating the number of pictures taken in a specific location (e.g., $photoTaken[loc : Location_type] = (int : Integer_type)$), and functional terms for recording the position of each actor (e.g., $at[prt : Participant] = (loc : Location_type)$) and robot (e.g., $atRobot[prt : Participant] = (loc : Location_type)$) in the area or for representing the updated situation of each location (e.g., $status[loc : Location_type] = (st : Status_type)$). Some terms may be used as *constant values*, and in this case the set of arguments taken as input by the single term is empty. For example, the term $generalBattery[] = (int : Integer_type)$ reflects the battery charge level stored in the power pack and used for recharging the battery of each robot (the term $batteryRecharging[] = (int : Integer_type)$ indicates the amount of battery that is charged after each recharging action). The terms $moveStep[] = (int : Integer_type)$ and $debrisStep[] = (int : Integer_type)$ reflects the amount of battery consumed respectively after the robot has been moved from a location to another one and after having removed debris from a specific location.

Finally, terms can also be used for expressing static relations over objects. The term $neigh[loc1 : Location_type, loc2 : Location_type] = (bool : Boolean_type)$ indicates all adjacent locations in the area (for example, $neigh[loc00, loc01] = true$), while the term $covered[loc : Location_type] = (bool : Boolean_type)$ reflects the locations covered by the network provided by the fixed antenna.

For each term, the process designer has to decide which ones are **relevant for adaptation** and which ones have not to be considered for that. A term that is considered as relevant for adaptation will be continuously monitored by the PMS, and if its value becomes different from the one expected after a task execution, the PMS will provide adaptation features.

⁵Predefined data types, like *Boolean_type* and *Integer_type*, can not be used as arguments of an atomic term.

⁶The special data type *Task* will be defined later and can be used only as argument of the pre-defined term *requires*.

Atomic Terms :Relevant for Adaptation :

```

at[prt:Participant] = (loc:Location_type)
evacuated[loc:Location_type] = (bool:Boolean_type)
status[loc:Location_type] = (st:Status_type)

```

Not Relevant for Adaptation :

```

atRobot[prt:Participant] = (loc:Location_type)
batteryLevel[prt:Participant] = (int:Integer_type)
photoTaken[loc:Location_type] = (int:Integer_type)
generalBattery[] = (int:Integer_type)
batteryRecharging[] = (int:Integer_type)
moveStep[] = (int:Integer_type)
debrisStep[] = (int:Integer_type)
neigh[loc1:Location_type,loc2:Location_type] = (bool:Boolean_type)
covered[loc:Location_type] = (bool:Boolean_type)

```

In addition to atomic terms, we allow the designer to define *complex terms*. They are declared as basic atomic terms, with the additional specification of a well-formed first-order formula φ that determines the truth value for the term. In our case study, we may need to express that an actor is connected to the network if s/he is in a covered location or if s/he is in a location adjacent to a location where a robot is located (and is thus connected through the robot):

```

isConnected(prt:Participant) {
  EXISTS(l1:Location_type, l2:Location_type, rbt:Participant).((at(act)=l1) AND
  (Covered(l1) OR (atRobot(rbt)=l2 AND Neigh(l1,l2) AND
  isRobotConnected(rbt))))).

```

The interpretation of complex terms derives from the corresponding first-order formula and is enacted at run-time by the IndiGolog interpreter. Formulae can be negated (NOT) and existentially or universally quantified (EXISTS and FORALL). A complex term can not appear in task effects and can not involve recursion. On the contrary, complex terms can be used within a task precondition. The description of the complex term *isRobotConnected* is given in the Appendix.

Concerning the definition of process tasks, the process designer is required to specify which tasks are applicable to the dynamic scenario under study. Those tasks will be stored in a specific *tasks repository*, and can be used for composing the control flow of the process (cf. Section 4.3.2) and for adaptation purposes.

Tasks Repository :

```

Tasks = {go, move, takephoto, evacuate, updatestatus, extinguishfire, chargebattery}

```

Each task is annotated with *preconditions* and *effects*. Preconditions are logical constraints defined as a conjunction of atomic terms, and they can be used to constrain the task assignment and must be satisfied before the task is applied, while effects establish the outcome of a task after its execution.

Definition 12. A task $t[x] \in \mathbb{T}$ is a tuple $t = (Act_t, x, Pre_t, Eff_t)$ that consists of:

- the name Act_t of the action involved in the enactment of the task (it often coincides with the task itself);
- a tuple of data objects x as input parameters;
- a set of preconditions Pre_t , represented as the conjunction of k atomic conditions defined over some specific terms, $Pre_t = \bigwedge_{l \in 1..k} pre_{t_l}$. Each pre_{t_l} can be represented as $\{v_j[y_j] \text{ op expr}\}$, where:
 - $v_j[y_j] \in \mathbf{V}$ is an atomic term, with $y_j \subseteq x$, i.e., admissible data objects for y_j need to be defined as task input parameters;
 - An **expr** can be a boolean value (if v_j is a boolean term); an input parameter identified by a data object (if v_j is a functional term); an integer number or an expression involving integer numbers and/or terms, combined with the arithmetic operators $\{+, -\}$ (if v_j is a integer term);
 - The condition **op** can be expressed as the equality ($==$) between boolean terms or functional terms and an admissible **expr**. On the contrary, if v_j is a integer term, it is possible to define the **op** condition as an expression that make use of relational binary comparison operators ($<, >, =, \leq, \geq$) and involve integer numbers and/or integer terms in the **expr** field.
- a set of deterministic effects Eff_t , represented as the conjunction of h atomic conditions defined over some specific terms, $Eff_t = \bigwedge_{l \in 1..h} eff_{t_l}$. Each eff_{t_l} (with $l \in 1..h$) can be represented as $\{v_j[y_j] \text{ op expr}\}$, where:
 - $v_j[y_j] \in \mathbf{V}$ and **expr** are defined as for preconditions.
 - The condition **op** may include assignment expressions to update the values of integer terms. A numeric effect consists of an assignment operator, the integer term to be updated and a integer number or a numeric expression. Assignment operators include (i) direct assignment ($=$), to assign to a integer term a value defined by an integer number; (ii) relative assignments, which can be used to increase ($+=$) or decrease ($-=$) the value of a integer term (additive assignment effects).

Note that if no preconditions are specified, then the task is always executable. Moreover, the process designer is required to make explicit if a task effect can be considered as *supposed* or *automatic*. A *supposed effect* indicates that the participant executing the task has to physically return an outcome for the supposed effect, that can be or not can be equal to the one declared during the task definition. If a supposed effect involves a relevant term, it is clear that the outcome returned assumes a great value for monitoring purposes. Otherwise, an effect can be flagged as *automatic*, meaning that when the task terminates, the effect is automatically applied without the need to consider the task outcomes.

As an example, the task *go* involves two input parameters *from* and *to* of type *Location_type*, representing the starting and arrival locations. The participant that will be assigned for executing the task *Go* is indicated with *PRT*. An instance of this task can be executed only if *PRT* is currently at the starting location *from* and provides the required capabilities for executing the task *go*. As a consequence of task execution, the actor moves from the starting to the arrival location, and this is reflected by assigning to the functional term $at[actor]$ the value *to* in the effect. The

fact that the effect of the task *go* is *supposed* means that after the execution of the task the participant will return the real outcome indicating her/his final position. Each task $t_i \in T$ together with its preconditions, effects and parameters can be represented as an XML annotation:

Description of the task *go* :

```
<task>
  <name>go</name>
  <parameters>
    <arg>from - Location_type</arg>
    <arg>to - Location_type</arg>
  </parameters>
  <precondition>at[PRT] == from AND isConnected[PRT] == true</precondition>
  <effects>
    <supposed>at[PRT] = to</supposed>
  </effects>
</task>
```

Let us observe now the task *move*. Like the task *go* it involves two input parameters *from* and *to* of type *Location_type*, that represent the starting and arrival locations given in input to the robot. We want to underline that the participant *PRT* which will be assigned to execute a *move* task will be for sure a robot. This can be inferred by analyzing the pre-defined terms *provides* and *requires*, that will be used by the PMS for understanding if a specific participant is capable to execute a specific task.

Description of the task *go* :

```
<task>
  <name>move</name>
  <parameters>
    <arg>from - Location_type</arg>
    <arg>to - Location_type</arg>
  </parameters>
  <precondition>atRobot[PRT] == from AND batteryLevel[PRT] >= moveStep[] AND
    isRobotConnected[PRT] == true
  </precondition>
  <effects>
    <supposed>atRobot[PRT] = to</supposed>
    <automatic>batteryLevel[PRT] -= moveStep[]</automatic>
  </effects>
</task>
```

An instance of the task *move* can be executed only if the robot *PRT* is currently at the starting location *from* and provides enough battery charge for executing the movement. As a consequence of task execution, the robot moves from the starting to the arrival location, and this is reflected by assigning to the functional term *atRobot[PRT]* the value *to* in the effect. This first effect of *move* has been flagged as *supposed*, meaning that after the execution of the task the robot will return the real outcome indicating its final position. However, since *atRobot* is not considered as a relevant term (see above), if the final position of the robot will differ with the one declared at design time, no adaptation is required. The second effect of *move* is *automatic*, and states that after the execution of a *move* task the battery level of the robot has to be decreased of a fixed quantity, corresponding to *moveStep[]*.

We can also represent exogenous events with SmartML. They reflect possible external events coming from the environment that can modify asynchronously atomic terms at run-time. In our case study, we can deal with three different exogenous events:

Exogenous Events :

```
Ex_events = {photoLost, fireRisk, rockSlide}
```

The definition of an exogenous event is similar to the classical task definition provided by SmartML. However, for defining an exogenous event there is no need to specify any precondition, and effects can only be considered as *automatic* (i.e., they are automatically applied to the involved terms when the exogenous event is caught). For example, the exogenous event *rockSlide(loc)* alerts about a rock slide collapsed in location *loc*, and its effect concerns to modify the value of the atomic term [*loc*] to the value '*debris*'.

Description of the exogenous event *rockSlide* :

```
<ex-event>
  <name>rockSlide</name>
  <parameters>
    <arg>loc - Location_type</arg>
  </parameters>
  <effects>
    <automatic>status[loc] = debris</automatic>
  </effects>
</ex-event>
```

Finally, we want to underline that the SmartML syntax allows to represent planning domains and problems with the complexity of those describable in PDDL version 2.2 [37], that is characterized for enabling the representation of realistic planning domains. The translation algorithms are described in Section 4.4.

4.3.2 Defining Processes in SmartPM through BPMN

Starting from a domain theory *D*, the control flow of a dynamic process in SmartPM can be defined through the BPMN notation. The Business Process Modeling Notation (BPMN) was released to the public in May 2004 by the BPMI Notation Working Group and was adopted as OMG standard⁷ for business process modeling in February 2006. BPMN provides a graphical notation for specifying business processes based on a flowcharting technique similar to activity diagrams from UML. In January 2001, the version 2.0 of the language has been released [7]. If compared to the previous specifications of the language, which provided only verbal descriptions of the graphic notations elements and modeling rules, BPMN 2.0 received a formal definition in the form of a meta-model, that defines the abstract syntax and semantics of the modeling constructs. The BPMN specification also provides a mapping between the graphic elements of the notation and the underlying constructs of BPEL (Business Process Execution Language) [95].

⁷<http://www.omg.org/spec/BPMN/2.0/>

Formally, in SmartPM we define a dynamic process as a directed graph consisting of tasks, gateways, events and transitions between them.

Definition 13. *Given a domain theory D and a set of tasks T , a dynamic process P is a tuple (N, L) where:*

- $N = T \cup E \cup W \cup X$ is a finite set of nodes, such that :
 - T is a set of task instances, i.e., occurrences of a specific task $t \in T$ in the range of the dynamic process;
 - E is a finite set of events, that consists of a single start event \circ and a single end event \odot ;
 - $W = W_{PS} \cup W_{PJ}$ is a finite set of parallel gateways, represented in the control flow with the \diamond shape with a “plus” marker inside.
 - $X = X_{ES} \cup X_{EJ}$ is a finite set of exclusive gateways, represented in the control flow with the \diamond shape with a “X” marker inside.
- $L = L_T \cup L_E \cup L_{W_{PS}} \cup L_{W_{PJ}} \cup L_{X_{ES}} \cup L_{X_{EJ}}$ is a finite set of transitions connecting events, task instances and gateways:
 - $L_T : T \rightarrow (T \cup W_{PS} \cup W_{PJ} \cup X_{ES} \cup X_{EJ} \cup \odot)$
 - $L_E : \circ \rightarrow (T \cup W_{PS} \cup X_{ES} \cup \odot)$
 - $L_{W_{PS}} : W_{PS} \rightarrow 2^T$
 - $L_{W_{PJ}} : W_{PJ} \rightarrow (T \cup W_{PS} \cup X_{ES} \cup \odot)$
 - $L_{X_{ES}} : X_{ES} \rightarrow 2^T$
 - $L_{X_{EJ}} : X_{EJ} \rightarrow (T \cup X_{ES} \cup W_{PS} \cup \odot)$

Note that the constructs used for defining a dynamic process are basically a subset of the ones definable through the BPMN notation. The intuitive meaning of these structures should be clear: an execution of the process starts at \circ and ends at \odot ; a *task* is an atomic activity executed by the process; *parallel splits* W_{PS} open parallel parts of the process, whereas *parallel joins* W_{PJ} re-unite parallel branches; exclusive gateways are used to create alternative flows in a process where only one of the path can be taken on the basis of a given condition. *Transitions* are binary relations describing in which order the flow objects (tasks, events and gateways) have to be performed, and determine the *control flow* of the dynamic process.

In the SmartPM system, we use the Eclipse BPMN editor⁸ for defining the control flow of a dynamic process. The BPMN editor provides visual, graphical editing and creation of BPMN 2.0 business processes. We have customized the tool for our needs, by adding the possibility to link domain theories written through SmartML to the control flow of the process.

Once a process is ready for being executed (i.e., the process designer has completed the definition of the domain theory and of the process control flow), the last step before sending the dynamic process to the IndiGolog engine for its enactment consists in instantiating the domain theory with a *starting condition*, which reflects different assignment of values to the atomic terms. We assume complete information about the starting condition, since the logical framework described in Chapter 3 works

⁸<http://eclipse.org/bpmn2-modeler/>

under the *closed-world assumption* [106]. Basically, this means we force the process designer to instantiate every atomic term with an admissible value that represent what is known in the starting state about the dynamic scenario. Specifically, the starting condition is a conjunction $\{v_1[y_1] == val_1 \wedge v_2[y_2] = val_2 \dots \wedge v_j[y_j] == val_j\}$, where val_j (with $j \in 1..m$) represents the j -th value assigned to the j -th atomic term.

4.3.3 The XML-to-IndiGolog Parser

From a technical point of view, a dynamic process built through Eclipse BPMN and annotated with SmartML is basically a file saved in the XML format that is taken as input from the XML-to-IndiGolog Parser component, which translates the process specification in situation calculus and IndiGolog readable formats, in order to make the process executable by the IndiGolog engine.

We show the rules used for translating the SmartML specification of a dynamic process in a SitCalc Theory:

- Each data type defined in SmartML, together with its associated data objects, is converted in a corresponding situation calculus predicate. For example, the pre-defined data type *Participant* corresponds to a predicate *Service* in situation calculus. The same is valid for the user-defined data types *Location_type* and *Status_type*, which are converted in the corresponding situation calculus versions, and for the all the tasks stored in the SmartML tasks repository. A situation calculus predicate *Task* indicates which tasks will be considered by the PMS for the process execution and adaptation (if required).
- The SmartML pre-defined terms, such as *provides* and *requires*, have again a corresponding translation as situation calculus predicates.
- Each atomic term will result in a situation calculus data fluent. For all those terms that are flagged as *relevant*, an expected fluent is also provided. For example, the *non relevant* term $atRobot[prt : Participant] = (loc : Location_type)$ will be translated in a situation calculus data fluent $atRobot_\varphi(prt, s)$, where prt is the participant (i.e., the robot) associated to the fluent and s is a situation term. A non relevant data fluent is not monitored for adaptation purposes. On the contrary, the *relevant* term $at[prt : Participant] = (loc : Location_type)$ will be converted in a data fluent $at_\varphi(prt, s)$ and in an expected fluent $at_\Psi(prt, s)$, that will respectively record the physical and expected positions of prt in every situation s , and are continuously monitored by the PMS.
- Tasks definitions are crucial for expressing situation calculus preconditions axioms and successor state axioms for each fluent. Let us consider, for example, the definition of the task *go*.

```

<task>
  <name>go</name>
  <parameters>
    <arg>from - Location_type</arg>
    <arg>to - Location_type</arg>
  </parameters>
  <precondition>at[PRT] == from AND
                isConnected[PRT] == true
  </precondition>

```

```

    <effects>
    <supposed>at[PRT] = to</supposed>
    </effects>
  </task>

```

First of all, the parser analyzes the structure of the XML element `<task>` for defining which admissible workitems including the task `go` can be executed by the IndiGolog PMS. By analyzing the structure of `go`, it is clear that it will require two inputs (specifically, two `Location_type` elements) and it will return only an outcome, that is again a `Location_type` element. This will led the parser to generate the following code:

$$\forall (from, to, id) \text{ s.t. } Identifier(id) \wedge Location_type(from) \wedge Location_type(to) \Rightarrow workitem(go, id, [from, to], [to]).$$

Then, the parser builds the precondition axiom for the task `go` by scanning the content of the XML element `<precondition>`:

$$\begin{aligned} Poss(assign(prt, id, go, [from, to], [to]), s) \Leftrightarrow \\ service(prt) \wedge workitem(go, id, [from, to], [to]) \wedge \\ At_{\varphi}(prt, s) = from \wedge isConnected(prt, s). \end{aligned}$$

Basically, the precondition axiom for the task `go` states that `go` can be assigned to a service `prt` in situation `s` if and only if `prt` is located in `from` and is connected to the network. Moreover, the precondition axiom verifies that by associating at run-time a whatever identifier to the task `go` with inputs `from` and `to` and expected effect equal to `to`, it is possible to define an admissible workitem.

Finally, the parser analyzes all the elements nested in the XML element `<effects>`. For each `supposed` effect, if the atomic term involved in the effect has been flagged as `relevant`, then the parser generates a data fluent for capturing the physical outcome of the task and an expected fluent for recording the expected value after task execution. In the case of the task `go`, the following fluents will be generated:

$$\begin{aligned} At_{\varphi}(prt, do(a, s)) = loc \equiv \\ (\exists loc \text{ s.t. } Location_type(loc) \wedge \\ a = release(prt, id, go, [from, to], [to], [loc])) \vee \\ (At_{\varphi}(prt, s) = loc \wedge \\ \neg \exists loc' \text{ s.t. } Location_type(loc') \wedge \\ a = release(prt, id, go, [from, to], [to], [loc']) \wedge (loc' \neq loc)). \end{aligned}$$

$$\begin{aligned} At_{\Psi}(prt, do(a, s)) = to \equiv \\ (\exists loc \text{ s.t. } Location_type(loc) \wedge \\ a = release(prt, id, go, [from, to], [to], [loc])) \vee \\ (At_{\Psi}(prt, s) = to \wedge \\ \neg \exists to', loc \text{ s.t. } Location_type(loc) \wedge Location_type(to') \wedge \\ a = release(prt, id, go, [from, to], [to'], [loc]) \wedge (to' \neq to)). \end{aligned}$$

When *prt* terminates the execution of the task *go*, its final position *loc* in situation *s* is stored in a specific data fluent At_{φ} , while the expected fluent At_{ψ} will assume the desired value *to* without considering the real outcome of the task.

If a task description includes *automatic effects* or *supposed effects* involving atomic terms flagged as *non relevant*, the parser will generate the successor state axioms only for the data fluents associated to the terms. For example, the definition of the task *move* (that instructs a robot to move from a position to another) involves these two kinds of effects (cf. Section 4.3.1), meaning that the parser will generate the successor state axioms for $AtRobot_{\varphi}$ (cf. Equation 3.19) and for $BatteryLevel_{\varphi}$ (cf. Equation 3.12).

- Each exogenous event formalized through SmartML will be translated in a corresponding situation calculus exogenous event. Specifically, the parser analyzes the effects provided by the execution of the exogenous event and applies it to the involved data fluent. Remember that exogenous events can only change values of data fluents, by leaving the expected reality untouched. Sometimes the effects of an exogenous event affect a data fluent for which there exists another task whose execution can modify the value of the same fluent. In such a case, the parser needs to combine the effects provided by an exogenous event with the effects returned after a task execution, by generating a single successor state axiom for the data fluent involved. For example, let us consider the SmartML definition of the task *takephoto* and of the exogenous event *photoLost*:

```

<task>
  <name>takephoto</name>
  <parameters>
    <arg>loc - Location_type</arg>
  </parameters>
  <precondition>at[PRT] == loc AND isConnected[PRT] == true</precondition>
  <effects>
    <supposed>photoTaken[loc] = true</supposed>
  </effects>
</task>

<ex-event>
  <name>photoLost</name>
  <parameters>
    <arg>loc - Location_type</arg>
  </parameters>
  <effects>
    <automatic>photoTaken[loc] = false</automatic>
  </effects>
</ex-event>

```

The task *takephoto* is used for instructing a service (specifically, a human actor) *PRT* to move in a location *loc* in order to take some pictures. When *PRT* terminates the execution of the task *takephoto*, the physical outcome of the task is stored in the atomic term $photoTaken[loc]$. The same happens when an exogenous event *photoLost* is captured by the PMS, but in this last case the atomic term is asynchronously switched to the *false* value. The parser, which

has already generated a data fluent $PhotoTaken_\varphi$, is in charge to associate a single successor state axiom that reflect every possibility of changing the fluent:

$$\begin{aligned}
 PhotoTaken_\varphi(loc, do(a, s)) = true \equiv & \\
 (\exists q_j \text{ s.t. } a = \text{release}(prt, id, \text{takephoto}, [loc], [true], [q_j]) \wedge (q_j = true)) \vee & \\
 (PhotoTaken_\varphi(loc, s) = true \wedge & \\
 ((a \neq \text{photoLost}(loc)) \vee & \\
 \neg \exists q'_j \text{ s.t. } (a = \text{release}(prt, id, \text{takephoto}, [loc], [true], [q'_j]) \wedge (q'_j = false))) &
 \end{aligned}$$

Consider also that the effect of $takephoto$ is a *supposed* one, and the involved term $photoTaken$ has been flagged as *relevant*. Therefore, the parser generates an expected fluent $PhotoTaken_\Psi$ where recording the desired outcome of $takephoto$.

$$\begin{aligned}
 PhotoTaken_\Psi(loc, do(a, s)) = true \equiv & \\
 \exists q_j \text{ s.t. } a = \text{release}(prt, id, \text{takephoto}, [loc], [true], [q_j]) \vee & \\
 (PhotoTaken_\Psi(loc, s) = true \wedge & \\
 (\neg \exists q'_j \text{ s.t. } (a = \text{release}(prt, id, \text{takephoto}, [loc], [false], [q'_j]))) &
 \end{aligned}$$

Clearly the successor state axiom of $PhotoTaken_\Psi$ is not affected by exogenous events, since they are though to modify only the physical reality.

- Complex terms are basically first-order formula evaluated on the atomic terms. They are directly translated in situation calculus abbreviations.
- The initial condition used for instantiating any value of atomic terms in the domain theory corresponds with the initial situation S_0 in IndiGolog.

The XML-to-IndiGolog Parser is also able to convert the control flow of a dynamic process, defined through the BPMN language, in a valid IndiGolog based process. The conversion is straightforward, and concerns to map every possible BPMN construct used for formalizing processes in SmartPM (cf. Section 4.3.2 for the list of admitted constructs) in a valid IndiGolog construct, as shown in Table 3.1. We omit here the full translation algorithm, but we refer the reader to [29] for a similar complete mapping between IndiGolog and WS-BPEL.

4.4 Building the Planning Domain and the Planning Problem

In order to exploit our planning-based recovery mechanism, every task/annotation/property associated to a dynamic process needs to be translated in PDDL. A PDDL definition consists of two parts: the domain and the problem definition. The planning domain is built when a dynamic process is ready to be executed, i.e., the process designer has completed the definition of the control flow of the dynamic process with an associated domain theory represented in the SmartML language (cf. Section 4.3.1). Specifically, the SmartML domain theory is given as input to a software module named *Domain Builder*, that is in charge to convert such specification in a planning domain that complies with the PDDL 2.2 [37] language, characterized for enabling the representation of realistic planning domains.

Basically, the Domain Builder starts analyzing the definition of atomic/complex terms and data types as formalized in the previous sections, and by making explicit the *actions* associated to each annotated task stored in the repository linked to the dynamic process under execution, together with the associated pre-conditions, effects and input parameters. Basically, the planning domain describes how terms may vary after a task execution, and reflects the contextual properties constraining the execution of tasks stored in a specific tasks repository.

In the following, we discuss how a domain theory defined through SmartML can be translated into a PDDL file representing the planning domain:

- each data type corresponds to an *object type* in the planning domain;
- boolean terms and complex terms have a straightforward representation as *relational predicates* and *derived predicates* (to model the dependency of given facts from other facts) in the planning domain;
- integer terms correspond to PDDL *numeric fluents*, and are used for modeling non-boolean resources (e.g., the battery level of a robot) in the planning domain;
- functional terms do not have a direct representation in PDDL 2.2, but may be replaced as relational predicates. Since an object function $f : Object^n \rightarrow Object$ map tuples of objects with domain types D^n to objects with co-domain type U , it may be coded in the planning domain as a relational predicate P of type (D^n, U) ;
- a given task, together with the associated pre-conditions and effects and input parameters, is translated in a PDDL *action schema*. An action schema describes how the relational predicates and/or numeric fluents may vary after the action execution. In the following, it is shown the PDDL representation of the task *go*:

```
(:action go
  :parameters (?x - service ?from - location_type ?to - location_type)
  :precondition (and (provides ?x movement) (free ?x)
                    (at ?x ?from) (isConnected ?x))
  :effect (and (not (at ?x ?from)) (at ?x ?to))
)
```

This task can be executed only if the participant denoted with x is free and is currently located in his/her starting location *from* and is connected to the network. The desired effect turns the value of the predicate $at(x, to)$ to *true* and $at(x, from)$ to *false*, meaning the actor moved in a new location. Let us note that respect the definition of the task *go* given with SmartML, the PDDL version of the task requires to explicitly specify in the preconditions the capabilities required by a generic service x for executing the specific task and the information about the availability of the service. Moreover, the list of arguments is augmented with the information related to the service x that will execute the task. These information are crucial for allowing the planner to schedule correctly the task during the building of the plan.

When something relevant happens during the process execution, a new planning problem is built at run-time on the same planning domain defined above, through

the description of an initial state and of a desired goal. As thoroughly discussed in Chapter 3, process adaptation is required when it is sensed a misalignment between $\Phi(s)$ (the “wrong” physical reality) and $\Psi(s)$ (the “safe” expected reality). The *Problem Builder* takes as input from the Synchronizer the two realities and converts them in a PDDL planning problem. Specifically:

- for each data type defined in the SitCalc Theory, all the possible object instances of that particular data type are explicitly instantiated as *constant symbols* in the planning problem (e.g., the fact that *act1*, *act2*, *act3*, *act4*, *rb1* and *rb2* are *Services*, *loc00*, ..., *loc33* are *Location_type*, etc.);
- Basically, the initial state of the planning problem corresponds to the physical reality that need to be adjusted, plus some information for helping the planner to build correctly the recovery procedure. The initial state is composed by the conjunction of all data fluents defined in the SitCalc Theory (that correspond to PDDL relational predicates or numeric fluents in the planning domain), and the information concerning which services are possibly free for executing the tasks in the recovery plan (i.e., if *s* is the situation where the deviation has been sensed, for each service *srv* we need to make explicit which fluent *Free(srv, s)* holds). Moreover, for each service *srv* and capability *c* defined in the SitCalc theory, it is required the list of the predicates *provides(srv, c)* equal to *true*.
- the *goal state* of the planning problem is a logical expression over facts. In our approach, the goal state is built in order to reflect a safe state to be reached after the execution of a recovery procedure. Therefore, it is composed by the conjunction of all the expected fluents (and their corresponding values) in situation *s*.

Appendix A shows the complete PDDL code that describes a planning domain and a planning problem for solving the exception depicted in Fig. 1.3 and 1.4.

Finally, we want to briefly introduce the role of the *Translator* component. Basically, it waits for a plan to be synthesized and, when the planner produces the recovery plan, the Translator converts it in a IndiGolog procedure that is executable by the IndiGolog PMS. The translation is straightforward, since the planner returns a sequence of workitems w_1, \dots, w_n (i.e., the tasks to be executed together with their input/outputs and the information about the services that will execute them), corresponding to an IndiGolog procedure $\delta_a = (\delta_{a_1}; \dots; \delta_{a_n})$. This procedure is sent to the Synchronizer component, that will build the adapted process and sends it back to the PMS.

4.5 SmartPM in Action

In this section, we show some screenshot of the SmartPM system while executing the process defined in our case study described in Section 1.3. In Fig. 4.3 it is shown the main window of the Eclipse BPMN modeler, that we use for building the control flow of our dynamic processes.

A Java module named the *SmartML Editor* (cf. Fig. 4.4(a)) is used for building domain theories based on the SmartML specification. The editor allows to load

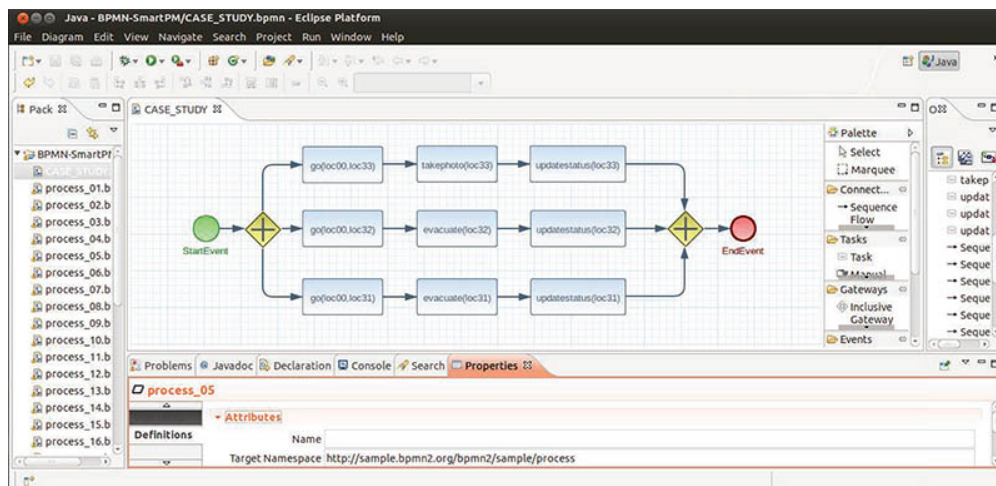


Figure 4.3. The Eclipse BPMN modeler.

previously saved specifications to be attached to a BPMN process, and to build new specifications customized on the basis of the designer needs. An *Edit* menu (cf. Fig. 4.4(b)) provides some guided masks for driving the user in the definition of the various objects of the specification. Basically, the editor drives the user inputs in order to facilitate the creation of a SmartML valid domain theory. The editor integrates also a *syntax checker* that verifies on the fly if the current specification is compliant with the SmartML formal one defined in Section 4.3.1.

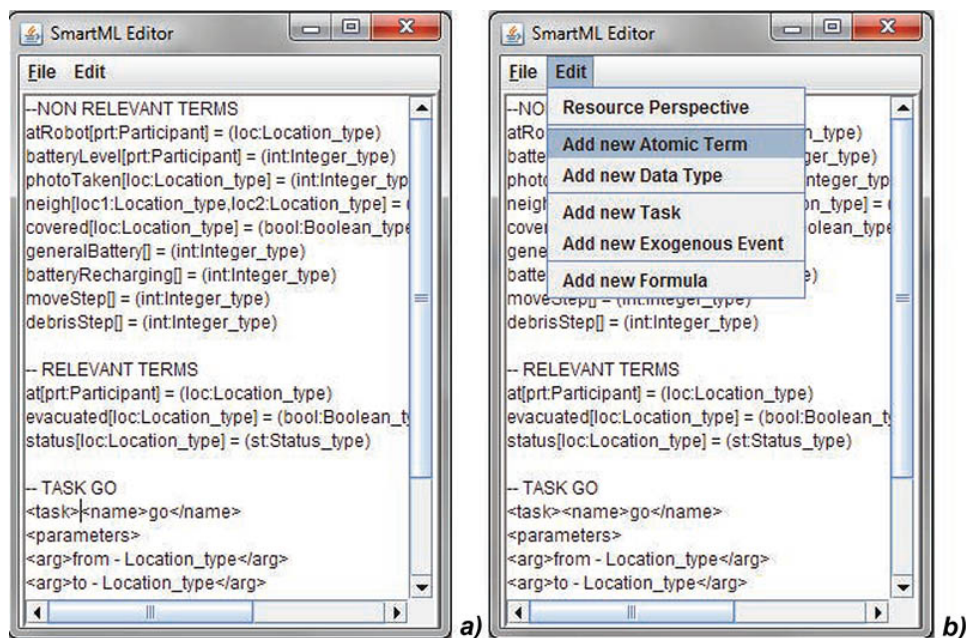


Figure 4.4. The SmartML editor (a) and the menu for inserting new terms/tasks/objects/formulae to the specification (b).

Once a process is ready for being executed, a new window is opened for allowing the process designer to instantiate the starting configuration for the atomic terms

in the domain theory and for binding actual input/output parameters with data objects provided with SmartML. The SmartML specification plus the BPMN process is now passed to the XML-to-IndiGolog parser component, that build an IndiGolog program and passes it to the IndiGolog PMS for the execution. In parallel, the Domain Builder component builds the PDDL planning domain to be used for a possible future adaptation of the process.

```

Terminal
File Edit View Search Terminal Help
,[loc32]])) * for execution
DEBUG 2: (EM) Start to execute the following action: 3,assign(act2,[workitem(g
o,id_4,[loc00,loc32],[loc32]]),javaswings(fb),assign(act2,[workitem(go,id_4,[lo
c00,loc32],[loc32]]))
DEBUG 2: (EM) Action * 3,assign(act2,[workitem(go,id_4,[loc00,loc32],[loc32]]))
,javaswings(fb),assign(act2,[workitem(go,id_4,[loc00,loc32],[loc32]])) * complet
ed with outcome: ok
>>>>>>>>>> ACTION EVENT:: Action * assign(act2,[workitem(go,id_4,[loc00,loc32]
,[loc32]])),3 * COMPLETED SUCCESSFULLY
DEBUG 2: Waiting for exogenous action to happen

=====> EXOGENOUS EVENT:: Exog. Action * readyToStart(act1,id_1,go) * occurre
d
***** PROGRAM:: Restarting step.
DEBUG 1: Sending nonsensing action * start(act1,id_1,go,[loc00,loc33],[loc33]) *
for execution
DEBUG 2: (EM) Start to execute the following action: 4,start(act1,id_1,go,[loc
00,loc33],[loc33]),javaswings(fb),start(act1,id_1,go,[loc00,loc33],[loc33])
DEBUG 2: (EM) Action * 4,start(act1,id_1,go,[loc00,loc33],[loc33]),javaswings(f
b),start(act1,id_1,go,[loc00,loc33],[loc33]) * completed with outcome: ok
>>>>>>>>>> ACTION EVENT:: Action * start(act1,id_1,go,[loc00,loc33],[loc33]),4
* COMPLETED SUCCESSFULLY
DEBUG 2: Waiting for exogenous action to happen

```

Figure 4.5. The main window of the IndiGolog PMS

Fig. 4.5 depicts the main window of the IndiGolog PMS showing the log of all actions exchanged between the PMS and services. In the screenshot, we can identify an *assign* action that assigns to service *act2* the task *go(id_4, [loc00, loc32], [loc32])*.

In Fig. 4.6 we show some screenshots representing our Task Handler for executing tasks. The Task Handler is implemented in Java, by the use of standard Java 2D graphical libraries. Specifically, when the IndiGolog PMS assigns a task to a service, this event is notified to the selected service through a popup window (cf. Fig. 4.6(a)). When a service is ready to start a task, it pushes the button *Start It*, and a *readyToStart* action is sent back to the IndiGolog PMS. In Fig. 4.5 we can also see the presence of a row `=====> EXOGENOUS EVENT`, that represents the fact that the PMS has captured the *readyToStart* action sent by the service. We have to underline that the actions sent by services to PMS (i.e., the *readyToStart* and the *finishedTask* actions) are recognized by the IndiGolog PMS as “good” exogenous events. In response, the IndiGolog PMS can command the service to start the task execution through a *start* action (cf. the bottom part of Fig. 4.5).

Now, let us suppose that service *act1* has been instructed to start the task *go(id_1, [loc00, loc33], [loc33])*. Service *act1* can choose one of the valid outcomes (i.e., a list of *Location_type* data objects) and pushes the button *End Task* when the


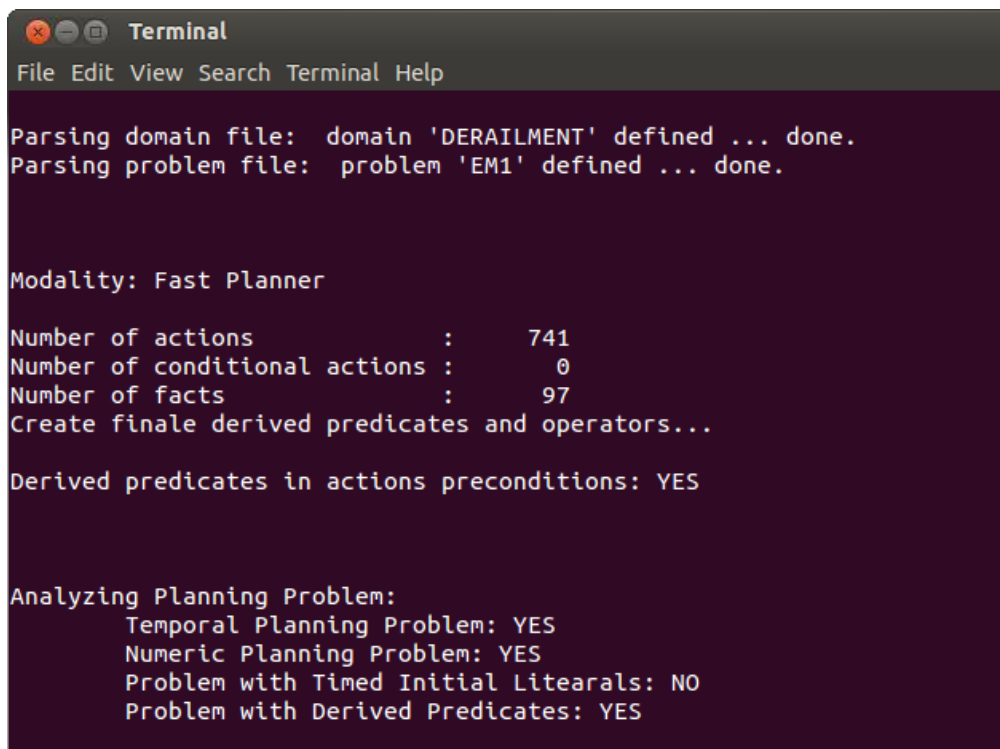


Figure 4.6. The Work-list Handler of SmartPM.

task is completed (cf. Fig. 4.6(b)). If the outcome provided by *act1* is different from the one expected (meaning it reaches a different location respect the one desired), the IndiGolog PMS senses the deviation, builds a planning problem that reflects the gap between physical and expected reality and launches the LPG-TD planner (cf. Fig. 4.7), which is in charge to synthesize the recovery plan. Note that the deviation we are analyzing is the same shown in our case study; i.e., we are supposing that *act1* has reached location *loc03* rather than location *loc33* (cf. Fig. 1.3(b)).

During the synthesis of the recovery plan, every running task is interrupted for the time required to the planner for building the plan and to the IndiGolog PMS for executing it (cf. Fig. 4.6(c)). When the planner finds a recovery procedure, it is passed back to the Synchronization component, that converts it in a executable IndiGolog process and sends it to the IndiGolog PMS for its enactment.



```

Terminal
File Edit View Search Terminal Help

Parsing domain file: domain 'DERAILMENT' defined ... done.
Parsing problem file: problem 'EM1' defined ... done.

Modality: Fast Planner

Number of actions          :      741
Number of conditional actions :        0
Number of facts           :        97
Create finale derived predicates and operators...

Derived predicates in actions preconditions: YES

Analyzing Planning Problem:
  Temporal Planning Problem: YES
  Numeric Planning Problem: YES
  Problem with Timed Initial Literals: NO
  Problem with Derived Predicates: YES

```

Figure 4.7. The main window of the LPG-td planner.

Since all the actors/robots need to be continually inter-connected to execute the process, the Planner finds a recovery procedure that first instructs the robots to move in specific positions for maintaining the network connection, and then re-assigns the task $go(loc03,loc33)$ to $act1$ (cf. Fig. 1.4). In Fig. 4.6(d) it is shown the assignment of a recovery task to the service $rb1$. Specifically, service $rb1$ executes the first task of the recovery procedure dealing with the deviation (cf. Fig. 1.4(a)).

We conclude this chapter by pointing out that the SmartPM System is a real working proof-of-concept prototype system that implements our approach to dynamic process adaptation shown in Section 3. While other approaches and systems rely on pre-defines rules to specify the exact behaviors when special events are triggered, here we simply model (a subset of) the running environment and the actions' effects, without considering any possible exceptional event. We argue that, in most of cases, modeling the environment, even in detail, is easier than modeling all possible exceptions.

Chapter 5

Validation

This chapter reports on performance evaluation and system validation activities. Specifically, in Section 5.1 we first report on experimental evaluation results, in terms of time needed for automatically adapting the dynamic process taken from our case study when exceptions of growing complexity arise. Then, in Section 5.2 we measure the effectiveness of SmartPM in finding recovery procedures by simulating the execution of thousands of processes instances having different structures of the control-flows.

5.1 Performances of SmartPM in Computing Recovery Procedures

In order to investigate the feasibility of the SmartPM approach, we performed some testing to learn the time amount needed for synthesizing a recovery plan for different adaptation problems. We made our tests by using the LPG-td planner¹ [45]. Such a planner is based on a stochastic local search in the space of particular “action graphs” derived from the planning problem specification. The basic search scheme of LPG-td is inspired to Walksat [118], an efficient procedure for solving SAT-problems. More details on the search algorithm and heuristics devised for this planner can be found at [45, 44].

We chose LPG-td as (i) it treats the full range of PDDL2.2³ [37] and (ii) even if it is primarily thought as a satisficing planner, it is able to compute also quality plans under a pre-specified metric. In fact, LPG-td has been developed in two versions: a version tailored to computation speed, named LPG-td.speed, which produces *sub-optimal plans*, and a version tailored for *plan quality*, named LPG-td.quality. LPG-td.speed generates sub-optimal solutions that do not prove any guarantee other than the correctness of the solution. LPG-td.quality differs from LPG-td.speed

¹LPG-td was awarded at the 4th International Planning Competition² (IPC 2004) as the “top performer in plan quality”.

³PDDLv2.2 [37] is characterized for enabling the representation of realistic planning domains, which include (in particular) actions and goals involving numerical expressions, operators with universally quantified effects or existentially quantified preconditions, operators with disjunctive or implicative preconditions, derived predicates and plan metrics. However, currently, our formalism does not allow to represent conditional and universally quantified effects.

Table 5.1. Time performances of LPG-td for adaptation problems of growing complexity.

Length of the recovery proc.	Problem instances	Avg. time needed for a sub-optimal sol. (sec)	Avg. length of a sub-optimal sol.	Avg. time needed for a quality sol. (sec)
1	29	6,769	3	7,768
2	36	7,213	3	16,865
3	32	7,846	4	24,123
4	25	8,128	5	37,017
5	21	8,598	8	39,484
6	17	8,736	9	52,421
7	13	9,188	13	73,526
8	12	9,953	14	81,414

basically for the fact that it does not stop when the first plan is found but continues until a stopping criterion is met. In our experiments, the optimization criteria was fixed as the minimum number of actions needed for the planner to reach the goal. Therefore, a quality plan uses the smallest number of actions needed for reaching a goal state. It is important to underline that satisficing planning is easy (polynomial), while optimal planning is hard (NP-complete) [54].

The experimental setup was performed with the test case shown in our running example. We stored in the task repository 20 different emergency management tasks, annotated with 28 relational predicates, 2 derived predicates and 4 numeric fluents, in order to make the planner search space very challenging. Then, we provided 185 different planning problems of different complexity, by manipulating ad-hoc the values of the initial state and the goal in order to devise adaptation problems of growing complexity.

As shown in Table 5.1, the column labeled as “Length of the recovery procedure” indicates the smallest number of actions needed for devising a plan of a specific length. Our purpose was to measure (in seconds) the computation time needed for finding a sub-optimal solution and a quality solution for problems that require a recovery procedure of growing complexity. The column labeled as “Average length of a sub-optimal solution” indicates the average number of actions that compose a sub-optimal solution for a problem of a given complexity. A sub-optimal solution is found in less time than a quality one, but generally it includes more tasks than the ones strictly needed. This means that when the complexity of the recovery procedure grows, the quality of a sub-optimal solution decreases. For example, as shown in table 5.1, on 21 different planning problems requiring a recovery procedure of length 5, the LPG-td planner is able to find, on average, a sub-optimal plan in 8,598 seconds (with 3 more tasks, on average) and a quality plan (which consists exactly of the 5 tasks needed for the recovery) in 39,484 seconds, without the need of any domain expert intervention. Consequently, the approach is feasible for medium-sized dynamic processes used in practice⁴.

5.2 Effectiveness of SmartPM in Adapting Processes

When developing a PMS with automatic adaptation features, an important aspect that needs to be analyzed concerns the *effectiveness* of the PMS in executing process

⁴We did our tests by using an Intel U7300 CPU 1.30GHz Dual Core, 4GB RAM machine.

instances that have different structures. We define effectiveness as the *ability of the PMS of executing every task included in the process control flow by adapting automatically the process itself if some failure arises, without the need of any manual intervention of the process designer at run-time.*

In order to measure the effectiveness of the SmartPM system in executing dynamic processes, we simulated the execution of 3600 process instances having different structure. Starting from tasks repositories with growing sizes, we generated process instances composed by sequences of tasks or by parallel branches. The simulation has concerned the testing of:

- 1200 process instances with a control flow composed by a sequence of tasks;
- 1200 process instances with a control flow composed by 3 parallel branches with tasks to be executed in concurrency;
- 1200 process instances with a control flow composed by 5 parallel branches with tasks to be executed in concurrency.

Our purpose was to measure the success rate in executing correctly the above process instances, by automatically adapting them if required. To this end, we developed a software module named the *SmartPM Simulator*, which is able to build automatically IndiGolog processes and SitCalc theories and simulating their execution on the basis of some customizable parameters:

- **Tasks repository size:** We allowed the generation of tasks repositories storing respectively 25, 50 or 75 tasks.
- **Number of available services:** This parameter indicates the number of services available in the initial situation for tasks assignment. We fixed this value to 5 services.
- **Maximum number of conditions in tasks preconditions/effects:** We allowed the generation of tasks having a maximum of 5 logical conditions in the precondition axioms, and a maximum of 5 different outcomes as effects.
- **Number of available fluents:** We allowed the generation of SitCalc theories composed by 50 relevant data fluents (meaning that, for each data fluent, the SmartPM Simulator automatically builds the corresponding expected fluent for monitoring possible tasks failures). The generated fluents can assume only boolean values.
- **Percentage of failures:** This parameter may assume two possible values (30% or 70%), and affects the percentage of tasks error during the process execution. For example, if a process instance is composed by 10 tasks, and the percentage of failures is equal to 70%, this means that 7 tasks will terminate with some outcomes different from the ones expected, and process adaptation is required.
- **Percentage of capabilities coverage:** This parameter can assume two possible values (30% or 70%), and affects the ability of each available service to execute the tasks stored in the repository. We assumed that each task is associated with a unique capability; therefore, if - for example - the tasks repository stores 75 tasks and the percentage of capabilities coverage is equal to 30%, this means that each available service will be able to execute at least

22 tasks in the repository. Moreover, we assume that for each task in the control flow there exists at least one service that is able to execute that task.

- **Number of tasks in the process control-flow:** Given a specific structure of the control flow, a fixed percentage of capabilities coverage and a fixed percentage of failures, for each possible size of the tasks repository the *SmartPM Simulator* generated 100 process instances with control flows composed respectively by 5, 8, 10, 12, 14, 15, 18, 20, 22, 25 tasks. Tasks composing the control flow are randomly picked from the tasks repository. For control flows with n concurrent branches, tasks are organized in a way that allows every branch to contain at least one task.

Test results are shown respectively in Figures 5.1, 5.2 and 5.3. Each figure reports a diagram related to the execution of 1200 process instances having (i) a control flow composed by a sequence of processes (cf. Fig. 5.1) (ii) a control flow composed by 3 parallel branches (cf. Fig. 5.2) and (iii) a control flow composed by 5 parallel branches (cf. Fig. 5.3). For each diagram, on the y-axis it is reported the size of the tasks repository, and on the x-axis is indicated the success rate (i.e., the effectiveness) in executing process instances with specific characteristics.

For example, let us consider the diagram in Fig. 5.1, that shows the effectiveness of *SmartPM* in executing processes with a structure composed by a sequence of tasks. When the size of the tasks repository is fixed to 25 tasks, the 4 colored bars denote respectively the effectiveness of the system in executing 400 process instances (each bar reflects the enactment of 100 instances) with a *percentage of capabilities coverage* and a *percentage of failures* varying from 30% to 70%. For example, the blue bar represents the effectiveness of the system in executing 100 process instances (with a variable dimension of the number of tasks composing the control flow, see above) with a percentage of failures of 30% and a percentage of capabilities coverage equal to 30%. According to the results shown in the diagram, the percentage of success is equal to 80%, meaning that 80 processes between the 100 that was executed have been correctly enacted and terminated, whereas for the other 20% the system did not find any recovery plan for dealing with the exception, meaning that a manual support to adaptation was required. The other bars, instead, indicate the effectiveness of the system by increasing the percentage of failures and the percentage of capabilities coverage. The analysis of the performed tests put in evidence some interesting aspects, that are valid for each of the 3 process structures we tested and for every possible size of the tasks repository:

- When the percentage of failures increases, the effectiveness of the *SmartPM* system decreases.
- Given a fixed percentage of failures, to increase the percentage of capabilities coverage means that the effectiveness of the *SmartPM* system increases, because there are more possibilities for a task in the repository to be selected by an available service.
- In general, to have a large tasks repository helps to increase the effectiveness of the *SmartPM* system, since the planner has more possibilities to select tasks for building a recovery procedure.

Moreover, we also compared the obtained data with the purpose to understand

if the structure of process instances may affect the effectiveness of SmartPM. To this end, for each possible combination of the percentage of failures with the percentage of capabilities coverage, we reorganized the collected data by generating 4 further diagrams that show the effectiveness of SmartPM after executing process instances whose control flows are composed by sequences of tasks or by parallel branches. The test results are shown respectively in Figures 5.4, 5.5 and 5.6 and 5.7. For example, in Fig. 5.5 we are comparing the execution of process instances with a percentage of failures equal to 30% and a percentage of capabilities coverage fixed to 70%. For instance, when the size of the tasks repository is equal to 75, the effectiveness of SmartPM in executing 100 process instances composed by a sequence of tasks is equal to 85%. The effectiveness decreases if the instances have tasks organized in 3 parallel branches (84%) and in 5 parallel branches (79%). In general, if the process control flow contains parallel branches, the effectiveness of the SmartPM system decreases as the number of parallel branches increases, since possibly more services are involved at the same time for tasks execution, by letting only few services available for process adaptation.

To sum, the execution of 3600 process instances with different structures was a valid test for measuring the effectiveness of SmartPM, that was able to complete correctly (i.e., without any intervention of domain experts at run-time) 2537 process instances, corresponding to an effectiveness of about 70,5%. Finally, we want to underline that the effectiveness of a PMS depends also by the ability of the process designer in formalizing the domain theory associated to a dynamic process. For example, by testing our case study (cf. Section 5.1) and by generating randomly an elevate number of possible failures, the planner was always able to find a proper recovery procedure and to adapt the process.




Figure 5.1. Effectiveness of SmartPM when executing processes composed by a sequence of tasks.




Figure 5.2. Effectiveness of SmartPM when executing processes composed by 3 parallel branches.




Figure 5.3. Effectiveness of SmartPM when executing processes composed by 5 parallel branches.

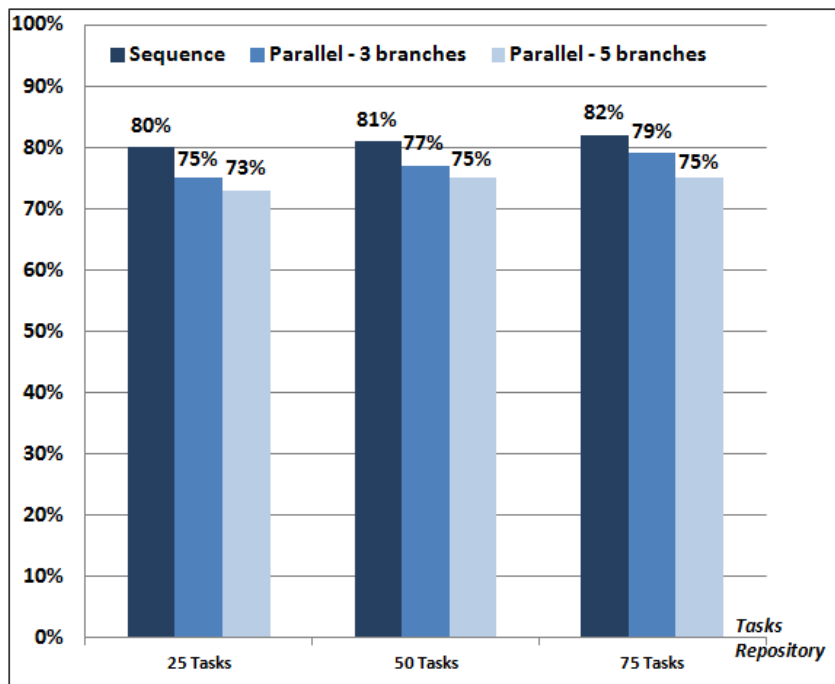


Figure 5.4. Measure of the SmartPM effectiveness by comparing process instances having different structures, with a percentage of capabilities coverage equal to 30% and a percentage of failures equal to 30%.




Figure 5.5. Measure of the SmartPM effectiveness by comparing process instances having different structures, with a percentage of capabilities coverage equal to 70% and a percentage of failures equal to 30%.




Figure 5.6. Measure of the SmartPM effectiveness by comparing process instances having different structures, with a percentage of capabilities coverage equal to 30% and a percentage of failures equal to 70%.




Figure 5.7. Measure of the SmartPM effectiveness by comparing process instances having different structures, with a percentage of capabilities coverage equal to 70% and a percentage of failures equal to 70%.

Chapter 6

Automatic Generation of Process Templates

Current workflow technology is based on the idea that, in general, there exists an underlying fixed process that can be used to automate the work. Once identified, a process is formalized into a *process model* and automated through a Process Management System that can execute it repeatedly. Conventional Business Process Management solutions require us to pre-define a detailed model of the process (control flow, data and resources) which captures every possible case (i.e., process instance) to be executed at run-time [142]. This approach works for processes where procedures are well known, repeatable (the work is done in the same way every time; i.e., there is enough similarity in executing each process instance) and can be planned in advance with some level of detail. As it emerges from the discussion conducted in Chapter 1, the need to deal with *knowledge-intensive processes* and *dynamic processes* and provide support for flexible process management has emerged as a leading research topic in the BPM domain [105].

In this chapter, we present an approach that allows us to automatically synthesize a *library of process templates* starting from a representation of the contextual domain in which the process is embedded in and from an extensive repertoire of tasks defined for such a context. A template depicts the best-practice procedure drawn up with whatever contextual information available at the time; it describes a recommended control flow for the process that does not only work in a specific state of the world, but can be enacted in a range of states satisfying the context conditions. In order to build process templates, we make use of *partial-order planning algorithms* (aka POP [140, 90]), which guarantee some interesting properties in the construction of the template:

- *Correctness*. Tasks composing the template are contextually selected from a specific repository and partially ordered in a way consistent with the context conditions to ensure that the template's objectives are achieved. Hence, a template is proven *correct* relative to the initial state of knowledge about the context.
- *Sound concurrency*. A process template has the property of *sound concurrency* in the execution of its activities. In fact, concurrent activities of a process template are proven to be effectively *independent* one from another (i.e.,

concurrent tasks cannot affect the same data). This means more flexibility during process execution. At runtime, the most appropriate execution path can be selected from those allowed by the design time process template definition, without the risk of interference between concurrent tasks.

- *Executability in partially known environments.* The use of classical plan-space algorithms for building the template requires complete knowledge of the starting state (i.e., it is not admitted that any fact is unknown). Once synthesized, a template can be executed in several different starting states, since it (usually) requires a fragment of the knowledge of the starting state to successfully achieve its objectives. We identify the *weakest preconditions* of process templates, and all the states satisfying such preconditions are good candidates for executing them. In this sense, templates can be enacted in *partially known environments*.

We exploit the idea behind POP of representing flexible plans that enables deferring decisions. Instead of committing prematurely to a complete, totally ordered sequence of actions, plans are represented as a partially ordered set, and only the required ordering decisions are recorded. A process template is generated on the basis of such a partially ordered set of activities, and we are able to identify what knowledge about the starting state is required for successful template execution. Moreover, we build step-by-step a library of process template specifications and support efficient retrieval of appropriate templates in partially known environments.

The rest of the chapter is organized as follows. Section 6.1 introduces a running example, derived from the main case study presented in Section 1.3. Section 6.2 discusses some preliminary notions around the use of partial-order planning techniques. Section 6.3 introduces the concept of process template. Section 6.4 presents the general approach used for synthesizing a library of process templates, starting from a planning domain and a planning problem. Section 6.5 gives some technical details about the algorithms used for computing process templates. Section 6.6 reports on experimental evaluation results. Section 6.7 discusses related work and Section 6.8 concludes by discussing benefits, limitations and future developments of the approach.

6.1 Case Study

Let us consider the emergency management scenario described in Fig. 6.1(a). It concerns a train derailment and depicts a map of the area (as a 4x4 grid of locations) where the disaster happened. For the sake of simplicity, we consider the same contextual information introduced in Section 1.3. Therefore, we have a derailed train composed of a locomotive (located in *loc33*) and two coaches (located in *loc32* and *loc31* respectively), and a response team sent to the derailment scene. The team is composed of four actors and two robots, initially located in *loc00*. If compared with the scenario introduced in Section 1.3, in this case we do not consider any information related to the network connection between actors and robots. We also remember that each process participant provides a set of specific capabilities. For example, actor *act1* is able to extinguish fire and take pictures, while *act2* and *act3* can evacuate people from train coaches. The two robots, instead, may remove debris

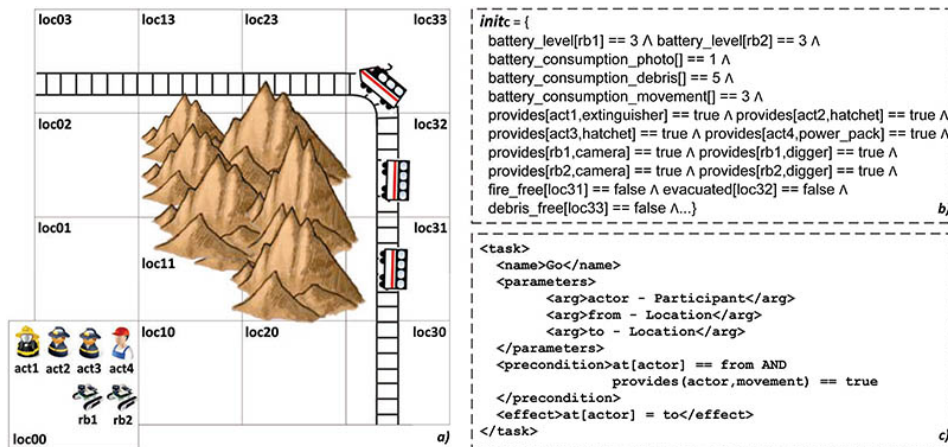


Figure 6.1. Area and context of the intervention.

from specific locations. Each robot has a battery and each action consumes a given amount of battery charge. When the battery of a robot is discharged, actor act_4 can charge it. Fig. 6.1(b) summarizes the above and shows the initial battery charge level of each robot.

Suppose now that the goal of an incident response plan defined for such a context is to evacuate people from the coach located in $loc32$, to extinguish fire in the coach in $loc31$ and finally to remove debris from $loc33$. If compared with the case study described in Section 1.3, now we do not provide any pre-built incident response plan for dealing with the disaster scenario. In fact, since the process may be different every time it is defined because it strictly depends on the actual contextual information (the positions of actors/robots, the location of every coach, the battery level of robots, etc.), it is unrealistic to assume that the process designer can pre-define all the possible process models for dealing with this environment. Moreover, if contextual data describing the environment are known, the synthesis of a process dealing with such environment is not straightforward, as the correctness of the process model is highly constrained by the values (or combination of values) of contextual data.

A simple approach to solving our problem is to build a process as a sequence of activities, e.g., the sequence of actions shown in Fig. 6.2. However, this solution is highly “inefficient”, since as many actions are independent, and they could be executed concurrently to reduce intervention time; e.g., a robot could removing debris in parallel with the extinguishing of the fire in $loc31$. But, at the same time, a process designer that has to design such a process may find difficult to organize activities for concurrent execution, since each action, for its executability, depends on the values of contextual data (e.g., a robot needs enough battery charge for moving into a location and removing debris). Also dependencies between actions play a key role in the definition of the process model (e.g., in order to evacuate people at $loc32$, a robot must have removed the debris beforehand). Finally, a process designer usually tends to represent more contextual information than that strictly needed for defining a process. For example, the execution of the process in Fig. 6.2 does not involve actor $act3$, meaning that any information concerning $act3$ (e.g., its capabilities, its location, etc.) is not required for synthesizing and




Figure 6.2. A process dealing with the scenario of Fig. 6.1(a).

executing the process. In order to overcome the above issues, we propose a solution that involves exploiting *partial-order planning* for generating a library of *process templates* for different contextual cases. Our templates provide *sound concurrency* in the execution of their activities and are executable in *partially known environments*.

6.2 Partial-Order Planning

Planning systems are problem-solving algorithms that operate on explicit representations of states and actions. The standard representation language for classical planners is known as the Planning Domain Definition Language (PDDL(cf. [46])); it allows one to formulate a problem PR through the description of the initial state of the world $init_{PR}$, the description of the desired goal condition $goal_{PR}$ and a set of possible actions. An action definition defines the conditions under which an action can be executed, called *preconditions*, and its *effects* on the state of the world. The set of all action definitions Ω is the *domain* PD of the planning problem. Each action $a \in \Omega$ has a precondition list and an effect list, denoted respectively as Pre_a and Eff_a . A planner that works on such inputs generates a sequence of actions (the *plan*) that corresponds to a path from the initial state to a state meeting the goal condition. In this chapter, we represent planning domains and planning problems making use of PDDL version 2.1¹(cf. [41]).

In the literature, there exists a wide range of different planning techniques, that are characterized by the specific assumptions made. In this chapter, we make use of *plan-space planning algorithms*. They differ from classical *state-space planning algorithms*, that explore only strictly linear sequences of actions directly connected to the start or goal, by devising totally ordered plans. A plan space is an implicit directed graph whose nodes are *partially specified plans* and whose edges correspond to refinement operations intended to further complete a partial plan, i.e., to achieve an open goal or to remove a possible inconsistency. In order to demonstrate our approach, we focus on Partial-Order Planning (POP) algorithms [90, 140], a specific type of plan-space planning algorithms. POP algorithms take as input a planning problem defined in PDDL and search the space of partial plans without committing to a totally ordered sequence of actions. They work back from the goal, by adding actions to the plan to achieve each subgoal. A tutorial introduction to POP algorithms can be found in [140].

¹PDDL 2.1 is characterized for enabling the representation of realistic planning domains, which include (in particular) actions with (linear) continuous numeric effects and effects dependent on the durations of the actions, actions and goals involving numerical expressions, operators with universally quantified effects or existentially quantified preconditions. However, currently, our formalism does not allow to handle negative preconditions, disjunctive preconditions and conditional effects.

Basically, a **partial plan** is a three-tuple $P = (A, O, CL)$, where $A \subseteq \Omega$ is a set of (ground) actions, O is a set of *ordering constraints* over A , and CL is a set of *causal links* over A . Ordering constraints O are of the form $a \prec b$, which is read as “ a before b ” and means that action a must be executed sometime before action b , but not necessarily immediately before. Causal links CL may be represented as $c \xrightarrow{p} d$, which is read as “ c achieves p for d ” and means that p is an effect of action c and a precondition for action d . It also asserts that p must remain *true* from the time of action c to the time of action d . In other words, the plan may not be extended by adding a new action that conflicts with the causal link and makes p false between c and d . Consequently, a precondition without a causal link requires further refinement to the plan to establish it, and is considered to be an *open condition* in the partial plan. Loosely speaking, the open conditions are preconditions of actions in the partial plan which have not yet been achieved in the current partial plan. More formally, an open condition is of the form (p, a) , where $p \in Pre_a$ and $a \in A$, and there is no causal link $b \xrightarrow{p} a$ (where b is any action of the partial plan P).

A classical POP algorithm starts with a null partial plan P and keeps refining it until a solution plan is found. The null partial plan contains two dummy actions $a_0 \prec a_\infty$ where the preconditions of a_∞ correspond to the top level goals $goal_{PR}$ of the problem, and the effects of a_0 correspond to the conditions in the initial state $init_{PR}$. Intuitively, a refinement operation avoids adding to the partial plan any constraints that are not strictly needed for addressing the refinement objective. This is called the *least commitment principle* [140]. The main advantage of the least-commitment philosophy is that decisions about action ordering are postponed until a decision is forced; constraints are not added to a partial plan unless strictly needed, thus guaranteeing flexibility in the execution of the plan and by possibly permitting actions to run concurrently. A **consistent** plan is defined as a plan with no cycles in the ordering constraints and no conflicts with the causal links. A consistent plan with no open conditions is a **solution** [90].

6.3 Process Templates

Our approach for the generation of a process template requires to explicitly model the contextual knowledge in which the dynamic process is embedded through some declarative rules (some pre-defined at design time, some known just before the synthesis of the template) and logical constraints expressed in terms of task preconditions and effects. Such information are given as input to an external partial-order planner [90, 140] that will be in charge to build a *process template*, i.e., a graph of activities reflecting the dynamic process required for solving the specific contextual problem. The language we are going to show can be seen as a dialect of SmartML (cf. Section 4.3.1), that we customized for describing process templates.

The synthesis of a dynamic process requires a tight integration of process activities and contextual data in which the process is embedded in. The context is represented in the form of a *Domain Theory* D , that involves capturing a set of tasks $t_i \in T$ (with $i \in 1..n$) and supporting information, such as the people/agents that may be involved in performing the process (roles or participants), the data and so forth. Tasks are collected in a specific repository, and each task can be considered as a single

step that consumes input data and produces output data. Data are represented through some ground atomic terms $v_1[y_1], v_2[y_2], \dots, v_m[y_m] \in \mathcal{V}$ that range over a set of tuples (i.e., unordered sets of zero or more attributes) y_1, y_2, \dots, y_m of *data objects*, defined over some *data types*. In short, a data object depicts an entity of interest. Under this representation, we consider possible values of a data type as constant symbols that univocally identify data objects in the scenario of interest.

Example 6.3.1. In our scenario we need to define data objects for representing participants (e.g., data type $Participant = \{act1, act2, act3, act4, rb1, rb2\}$), capabilities (e.g., data type $Capability = \{extinguisher, movement, \dots hatchet\}$) and locations in the area (e.g., data type $Location = \{loc00, loc10, \dots loc33\}$).

Each tuple y_j may contain one or more data objects belonging to different data types. The domain $dom(v_j[y_j])$ over which a term is interpreted can be of various types: (i) *Boolean*: $dom(v_j[y_j]) = \{true, false\}$, (ii) *Integer*: $dom(v_j[y_j]) = \mathbb{Z}$, (iii) *Functional*: the domain contains a fixed number of data objects of a designated type. Terms can be used to express properties of data objects (and relations over objects) and argument types of a term (taken from the set of data types previously defined) represent the finite domains over which the term is interpreted.

Example 6.3.2. In our example, we may need boolean terms for expressing the presence of a fire in a location (e.g., $fire_free[loc : Location] = (bool : Boolean)$), integer terms for representing the battery charge level of each robot (e.g., $battery_level[prt : Participant] \in \mathbb{Z}$) or functional terms for recording the position of each actor in the area (e.g., $at[prt : Participant] = (loc : Location)$).

Moreover, since each task has to be assigned to a participant that provides all of the skills required for executing that task, there is the need to consider the participants “capabilities”. This can be done through a boolean term $provides[prt : Participant, cap : Capability]$ that is *true* if the capability cap is provided by prt and *false* otherwise.

Each task is annotated with *preconditions* and *effects*. Preconditions are logical constraints defined as a conjunction of atomic terms, and they can be used to constrain the task assignment and must be satisfied before the task is applied, while effects establish the outcome of a task after its execution. Note that, as shown in Fig. 6.3(a), our approach treats each task as a “black box” and no assumption is made about its internal behavior (we consider the task execution as an instantaneous activity).

Definition 14. A task $t[x] \in \mathbb{T}$ consists of:

- the name of the action involved in the enactment of the task (it often coincides with the task itself);
- a tuple of data objects x as input parameters;
- a set of preconditions Pre_t , represented as the conjunction of k atomic conditions defined over some specific terms, $Pre_t = \bigwedge_{l \in 1..k} pre_{t_l}$. Each pre_{t_l} can be represented as $\{v_j[y_j] \text{ op } expr\}$, where:
 - $v_j[y_j] \in \mathcal{V}$ is an atomic term, with $y_j \subseteq x$, i.e., admissible data objects for y_j need to be defined as task input parameters;

- An **expr** can be a boolean value (if v_j is a boolean term); an input parameter identified by a data object (if v_j is a functional term); an integer number or an expression involving integer numbers and/or terms, combined with the arithmetic operators $\{+, -\}$ (if v_j is a integer term);
- **op** $\in \{<, >, ==, \leq, \geq\}$ is a relational operator. The condition **op** can be expressed as the equality ($==$) between boolean terms or functional terms and an admissible **expr**. On the contrary, if v_j is a integer term, it is possible to define the **op** condition as an expression that make use of relational binary comparison operators ($<, >, =, \leq, \geq$) and involve integer numbers and/or integer terms in the **expr** field.
- a set of deterministic effects Eff_t , represented as the conjunction of h atomic conditions defined over some specific terms, $Eff_t = \bigwedge_{l \in 1..h} eff_{t_l}$. Each eff_{t_l} (with $l \in 1..h$) can be represented as $\{v_j[y_j] \text{ op expr}\}$, where:
 - $v_j[y_j] \in \mathbf{V}$ and **expr** are defined as for preconditions.
 - **op** $\in \{=, +=, -=\}$ is used for assigning ($=$) to a term a value consistent with the **expr** field or for incrementing ($+=$) or decrementing ($-=$) an integer term by that value.

Note that if no preconditions are specified, then the task is always executable. As we will see in Section 6.4, the use of classical partial-order planning techniques for synthesizing process templates imposes some limitation in the expressiveness of the language used for defining the Domain Theory D. Specifically, negative preconditions are not admitted (e.g., the use of the NOT operator is forbidden and all the atomic conditions that require to evaluate if a boolean term is equal to *false* will be ignored) and we assume that all effects are deterministic.

Example 6.3.3. The task *Go* (cf. Fig. 6.1(c)) involves two input parameters *from* and *to* of type *Location*, representing the starting and arrival locations, and an input parameter *actor* of type *Participant* representing the first responder that will execute the task. An instance of this task can be executed only if *actor* is currently at the starting location *from* and provides the required capabilities for executing the task *Go*. As a consequence of task execution, the actor moves from the starting to the arrival location, and this is reflected by assigning to the functional term $at[actor]$ the value *to* in the effect.

Modeling a business process involves representing how a business pursues its objectives/goals. Objectives are represented in terms of a process goal to be satisfied. The goal may vary depending on the specific *Process Case C* to be handled. A case C reflects an instantiation of the domain theory D with a starting condition $init_C$ and a goal condition $goal_C$. Both conditions are conjunctions of ground atomic terms. We do not assume complete information about the starting condition; this means we allow a process designer to instantiate only the ground atomic terms necessary for representing what is known about the starting state, i.e., $init_C = \{v_1[y_1] == val_1 \wedge \dots \wedge v_j[y_j] == val_j\}$, where val_j (with $j \in 1..m$) represents the j -th value assigned to the j -th atomic term. Fig. 6.1(b) shows a portion of $init_C$ concerning the scenario depicted in Fig. 6.1(a). The goal is a condition represented as a conjunction of some specific terms we want to make true through the execution of the process.




Figure 6.3. Task anatomy (a), causality (b) and concurrency (c) in a process model.

Example 6.3.4. For example, in the scenario shown in Section 6.1, the goal has to be represented as $: goal_C = \{fire_free[loc31] == true \wedge evacuated[loc32] == true \wedge debris_free[loc33] == true\}$.

The syntax of goal conditions is the same as for tasks preconditions. A state is a complete assignment of values to atomic terms in V . Given a case C , an intermediate state $state_{C_i}$ is the result of i tasks performed so far, and atomic terms in V may be thought of as “properties” of the world whose values may vary across states.

Definition 15. A task t can be performed in a given state $state_{C_i}$ (and in this case we say that t is **executable** in $state_{C_i}$) iff $state_{C_i} \vdash Pre_t$, i.e. $state_{C_i}$ **satisfies** the preconditions Pre_t for the task t .

Moreover, if executed, the effects Eff_t of t modify some atomic terms in V and change $state_{C_i}$ into a new state $state_{C_{i+1}} = update(state_{C_i}, Eff_t)$. The *update* function returns the new state obtained by applying effects Eff_t on the current state $state_{C_i}$. Starting from a domain theory D , a *Process Template* captures a partially ordered set of tasks, whose successful execution (i.e., without exceptions) leads from $init_C$ to $goal_C$. Formally, we define a template as a directed graph consisting of tasks, gateways, events and transitions between them.

Definition 16. Given a domain theory D , a set of tasks T and a case C , a *Process Template* PT is a tuple (N, L) where:

- $N = T \cup E \cup W$ is a finite set of nodes, such that :
 - T is a set of tasks instances, i.e., occurrences of a specific task $t \in T$ in the range of the process template;
 - E is a finite set of events, that consists of a single start event \circ and a single end event \odot ;
 - $W = W_{PS} \cup W_{PJ}$ is a finite set of parallel gateways, represented in the control flow with the \diamond shape with a “plus” marker inside.
- $L = L_T \cup L_E \cup L_{W_{PS}} \cup L_{W_{PJ}}$ is a finite set of transitions connecting events, task instances and gateways:
 - $L_T : T \rightarrow (T \cup W_{PS} \cup W_{PJ} \cup \odot)$
 - $L_E : \circ \rightarrow (T \cup W_{PS} \cup \odot)$
 - $L_{W_{PS}} : W_{PS} \rightarrow 2^T$
 - $L_{W_{PJ}} : W_{PJ} \rightarrow (T \cup W_{PS} \cup \odot)$

Note that the constructs used for defining a template are basically a subset of the ones definable through the BPMN notation. Intuitively, an execution of the process starts at \circ and ends at \odot ; a *task* is an atomic activity executed by the process; *parallel splits* W_{PS} open parallel parts of the process, whereas *parallel joins* W_{PJ} re-unite parallel branches. *Transitions* are binary relations describing in which order the flow objects (tasks, events and gateways) have to be performed, and

determine the *control flow* of the template. A transition $l \in L$ is usually represented as $p \rightarrow q$, where $(p, q) \in N$. This represents the fact that there is a transition from the flow object p to the flow object q . For $n \in N$, $IN(n)/OUT(n)$ denotes the set of incoming/outgoing transitions of n , with the following restrictions :

- Only one outgoing/incoming flow may be associated with \circ and \odot respectively, i.e., $IN(\circ) = 0$, $OUT(\circ) = 1$, $IN(\odot) = 1$, $OUT(\odot) = 0$
- Each parallel split $w_{PS} \in W_{PS}$ accepts one incoming flow and more outgoing flows, i.e., $IN(W_{PS}) = 1$, $OUT(W_{PS}) > 1$
- Each parallel join $w_{PJ} \in W_{PJ}$ accepts more incoming flows and one outgoing flow, i.e., $IN(W_{PJ}) > 1$, $OUT(W_{PJ}) = 1$;
- Every task $t \in T$ is connected exactly to one incoming/outgoing flow, i.e., $IN(t) = 1$, $OUT(t) = 1$.

For example, in Fig. 6.3(b) we have a relation of *causality* between tasks t_a and t_b , stating that t_a must take place before t_b happens as t_a achieves some of t_b 's preconditions. An important feature provided by a process template is *concurrency*, i.e., several tasks can occur concurrently. In Fig. 6.3(c) an example of concurrency between t_1 and t_2 is shown. In order to represent two or more concurrent tasks in a template, the process designer makes use of the parallel gateways, that indicate points of the template in which tasks can be carried out concurrently. A parallel gateway may act as a *divergence element* (parallel split W_{PS}) or *convergent element* (parallel join W_{PJ}). As a point of divergence, the diamond shape is used when many tasks have to be carried out at the same time and in any order, which indicates that all transitions that exit this shape will be enabled together. As a point of convergence, the diamond shape is used to synchronize paths that exit a divergence element. This means that a process template is a *graph of tasks* (i.e., not a sequence) that imposes a *partial order* on their execution. A *linearization* of a process template is any linear ordering of the tasks that is consistent with the ordering constraints of the template itself [47]; i.e., a linearization of a partial order is a potential *execution path* of the template from the start event \circ to the end event \odot . For example, the template in Fig. 6.3(c) has two possible execution paths $r_1 = [\circ; t_a; t_1; t_2; t_b; \odot]$ and $r_2 = [\circ; t_a; t_2; t_1; t_b; \odot]$.

Definition 17. Given a process template PT and an initial state $state_{C_0} \vdash init_C$, a state $state_{C_i}$ is said to be **reachable** with respect to PT iff there exists an execution path $r = [\circ; t_1; t_2; \dots; t_k; \odot]$ of PT and a task t_i (with $i \in 1..k$) such that $state_{C_i} = update(update(\dots update(state_{C_0}, Eff_{t_1}) \dots, Eff_{t_{i-1}}), Eff_{t_i})$.

Definition 18. A task t_1 **affects** the execution of a task t_2 ($t_1 \triangleright t_2$) iff there exists a reachable state $state_{C_i}$ of PT (for some initial state $state_{C_0}$) such that:

- $state_{C_i} \vdash Pre_{t_2}$
- $update(state_{C_i}, Eff_{t_1}) \not\vdash Pre_{t_2}$

This means that Eff_{t_1} modify some terms in \mathbf{V} that are required as preconditions for making t_2 executable in $state_{C_i}$.

Definition 19. Given a process template PT , a case C and an initial state $state_{C_0} \vdash init_C$, an execution path $r = [\circ; t_1; t_2; \dots; t_k; \odot]$ (where $k = |T|$) of PT is said to be **executable** in C iff:

- $state_{C_0} \vdash Pre_{t_1}$
- for $1 \leq i \leq k - 1$, $update(state_{C_{i-1}}, Eff_{t_i}) \vdash Pre_{t_{i+1}}$
- $update(state_{C_{k-1}}, Eff_{t_k}) = state_{C_k} \vdash goal_C$

Definition 20. A process template PT is said to be **executable** in a case C iff any execution path of PT is executable in C .

The concept of execution path of a template helps in defining formally the *independence* property between concurrent tasks:

Definition 21. Given a process template PT , a task t_x is said to be **concurrent** with a task t_z iff there exist two execution paths r_1 and r_2 of PT such that $r_1 = [\circ; t_1; t_2; \dots; t_x; \dots; t_z; \dots; \odot]$ and $r_2 = [\circ; t_1; t_2; \dots; t_z; \dots; t_x; \dots; \odot]$.

Definition 22. Two concurrent tasks t_1 and t_2 are said to be **independent** ($t_1 \parallel t_2$) iff $t_1 \not\triangleright t_2$ and $t_2 \not\triangleright t_1$; that is, t_1 does not affect t_2 and vice versa.

6.4 On Synthesizing a Library of Process Templates

Our approach is focussed on the development and use of a *library* of *process templates*. These are reusable processes that achieve specified goals of interest in a range of starting states, i.e., any starting state that satisfies the template's required preconditions. We claim that in many cases, process template development can be partially automated through the use of AI planning tools.

Specifically, we focus on the use of a POP-based tool that can synthesize complex concurrent process models that are hard for humans to develop correctly (it is difficult for human designers to ensure that concurrent tasks never interfere with each other). The process designer's role is to specify the domain and context in which the template may be executed. Our POP-based tool can then be used to synthesize some candidate process models for the template. If the tool fails to generate a process model or the generated processes are of insufficient quality (e.g., they are too time consuming, unreliable, or lack concurrency), the designer can refine the domain theory and case to obtain better solutions. Once a satisfactory template has been obtained, it is added to the library. The POP-based tool automatically identifies the required preconditions for the template to achieve its goal, meaning the template can be reused whenever a case that matches the template's preconditions arises.

In a dynamic process domain, there is a wide range of cases/contexts to handle. New cases often arise and the requirements for the system frequently evolve. The designer develops and maintains the template library over time, in order to have templates that handle effectively most the cases that arise. The library also stores the templates specifications, i.e., their process domains, goals, and initial conditions/cases. New cases are often variants of existing cases and the designer will be able to adapt existing domain and case specifications to generate templates for the new cases using the tool. In the following, we describe an architecture and methodology for developing such a library-based approach.

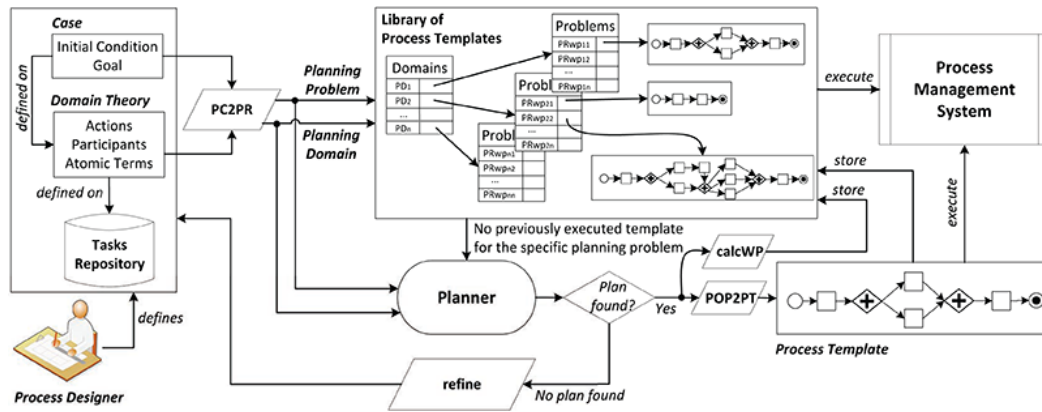


Figure 6.4. Overview of the general approach.

6.4.1 The General Framework

Our approach to the definition of a process template (cf. Fig. 6.4) requires a fundamental shift in how one thinks about modeling business processes. Instead of defining a process model “by hand”, here, the process designer has to address his/her efforts to specifying the Domain Theory D and the Case C to be handled. In particular, the process designer has to “guess” the starting condition $init_C$, by instantiating only some atomic terms, that ones needed for depicting the context the user has in mind. This means that $init_C$ can be partially specified, i.e. not all terms need to be instantiated with some value. Also the goal condition $goal_C$ is required, since it reflects the target state after having executed the template.

Example 6.4.1. Let us consider the scenario depicted in Section 6.1, represented with a Domain Theory D_1 and a goal condition $goal_{C_1} = \{fire_free[loc31] == true \wedge evacuated[loc32] == true \wedge debris_free[loc33] == true\}$. Since the process designer may be interested in an emergency process that involves the fewest participants, s/he can start by modeling a starting condition $init_{C_1}$ with information involving only actors $act1$ and $act2$ and the robot $rb1$, while terms involving $act3$, $act4$ and $rb2$ are not explicitly instantiated in $init_{C_1}$.

A specific module named PC2PR is in charge of converting the Domain Theory D and the Case C just defined into the corresponding Planning Domain PD and Planning Problem PR specified in PDDL version 2.1 (cf. [41]). Basically, PC2PR implements a function $f_{PC2PR} : (D, init_C, goal_C) \rightarrow (PD, init_{PR}, goal_{PR})$. Since the use of classical partial-order algorithms for synthesizing the template requires the initial state of PR to be a complete state, we make the closed world assumption [106] and assume that every atomic term $v_j[y_j]$ that is not explicitly specified in $init_C$ is assumed to be false (if $v_j[y_j]$ is a boolean term) or “not assigned” (if $v_j[y_j]$ is an integer or a functional term) in $init_{PR}$. Technical details of the algorithm employed in PC2PR are shown in Section 6.5.1.

At the heart of our approach lies a library of process templates built for specific planning domains and problems/cases. If library templates exist for the current values of PD and PR , we can retrieve an appropriate template and allow to execute it through an external PMS. However, if no template exists for the current values of PD

and PR, we can invoke an external POP planner on these same inputs. The planner will try to synthesize a plan fulfilling the goal condition $goal_{PR}$. If the planner is unable to find a plan, this suggests there are some missing elements in the definition of the Domain Theory D or in the Case C. Hence, to address this particular case, one can try to *refine* the case C and add information so that it becomes possible to generate a plan. There are many ways to strengthen a problem description, such as adding to the starting condition $init_C$ some terms initially ignored (e.g., to specify the position of every participant), or adding new objects in D or new activities in T (e.g., if a task for extinguish fire is missing). Our approach assumes that one specifies the context step-by-step, and requires the process designer to contribute to the system.

Example 6.4.2. If the planner is invoked with $init_{PR_1}$ (devised by applying f_{PC2PR} on the triple $D_1, init_{C_1}, goal_{C_1}$), it will not be able to find any plan for the specific problem. This is because *rb1* does not have enough battery charge for moving and removing debris. The designer can try to add new information to the problem description by instantiating in $init_{C_1}$ all those atomic terms related to actor *act4*, the only one able to charge robot batteries, and devises a new starting condition $init_{C_2}$ (and, consequently, a new initial planning state $init_{PR_2}$). A planner invoked with $init_{PR_2}$ is finally able to find a consistent plan P_1 satisfying $goal_{PR_1}$.

When the POP planner is able to find a partially ordered plan P consistent with the actual contextual information, three further steps are required. First of all, there is the need to translate the plan just found into a template PT that preserves the ordering constraints imposed by the plan. A **solution plan** is a three-tuple $P = (A, O, CL)$ that specifies the causal relationships for the actions $a_i \in A$, but without specifying an exact order for executing them. Since the set of actions $A \in P$ and the set of ordering constraints O over A must be explicitly expressed as nodes and transitions for the template's control flow (as well as their intrinsic ordering), we developed a module POP2PT implementing a function $f_{POP2PT} : P \rightarrow PT$ that takes as input P and converts it into a template PT. A detailed description of the module implementing POP2PT is presented in Section 6.5.2.

Example 6.4.3. By applying f_{POP2PT} to P_1 , we devise the template PT_1 in Fig. 6.5(a). Dashed arrows are causal links that imply an ordering constraint between pairs of tasks. For example, the ordering constraint between $Go[act1,loc00,loc31]$ and $ExtinguishFire[act1,loc31]$ is derived from the fact that *Go* has the effect $at[act1]=loc31$ that is needed by *ExtinguishFire* as precondition (i.e., *act1* has to be located in *loc31* for extinguish the fire in those location).

Secondly, our approach aims to *infer* the weakest preconditions w_{PT} under which the process template will achieve its objectives, i.e., to identify the least amount of information about the starting state that is required for the template to achieve its goal. The module we use for inferring w_{PT} is called **calcWP** and works by analyzing the set of causal links CL computed by the POP planner, to see which logical facts f_k are involved in causal links that originate from the dummy start action a_0 and end in some $a_k \in A$. More formally:

$$\forall (cl_k, f_k, a_k) \text{ s.t. } cl_k = (a_0 \xrightarrow{f_k} a_k) \in CL, \text{ then } f_k \in w_{PT}. \quad (6.1)$$

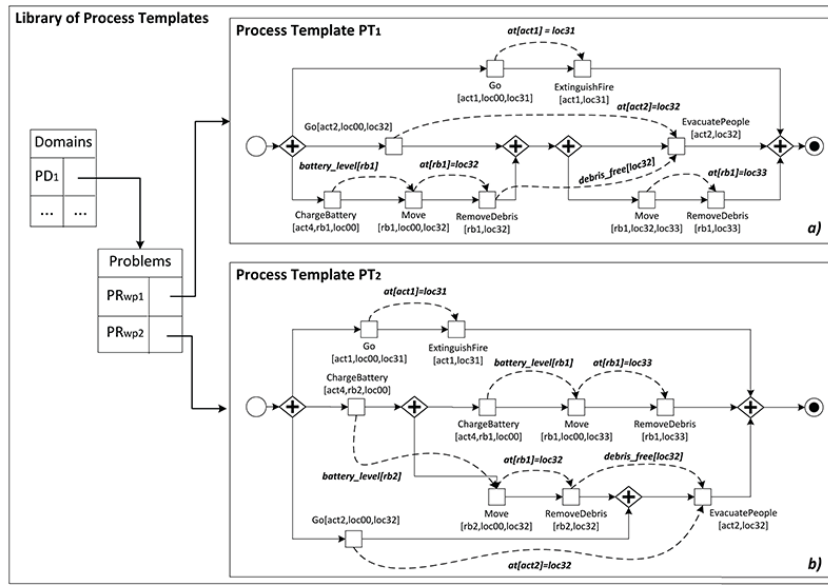


Figure 6.5. Templates dealing with the scenario in Fig. 6.1.

This exploits the fact that the effects of $a_0 \in A$ reflect the “true facts” in the starting state $init_{PR}$ that are actually used by the plan and ensure that it is executable and achieves the goal. A first result is that all atomic terms not explicitly instantiated in $init_C$ will correspond to facts assumed to be *false* in $init_{PR}$. As the closed world assumption suggests, false facts do not appear in $init_{PR}$ and are not eligible for being used as effects of a_0 . Secondly, by going through causal links as specified by (6.1), it is possible to identify all logical facts that - although being instantiated in $init_{PR}$ - are not required by the plan P . Basically, w_{PT} is the conjunction of those facts strictly required for executing the plan P (and, consequently, the devised template PT), and is used for devising a new problem $PR_{wp} = \{w_{PT}, goal_{PR}\}$.

Example 6.4.4. If we invoke $calcWP$ on the causal links devised from P_1 , we may easily infer w_{PT_1} . This means that for executing PT_1 (cf. Fig. 6.5(a)) we simply need to know the positions and capabilities of actors $act1$, $act2$, $act4$ and robot $rb1$; the other contextual information is not strictly needed for a correct execution of the template.

Thirdly, after the process template PT has been synthesized starting from P , it can be stored in our library together with information about the planning domain PD and abstracted problem PR_{wp} . Specifically, for every different planning domain PD devised through our approach, there is a pointer to a list of different abstracted planning problems PR_{wp} used for obtaining consistent plans in previous executions of our tool, together with the devised process templates. When a process designer defines a new Domain Theory D_{new} and a Case C_{new} , the system checks if the corresponding planning domain PD_{new} and problem PR_{new} (obtained by applying f_{PC2PR} to D_{new} and C_{new}) are already present in our library. If the library contains a planning domain PD and an abstracted planning problem PR_{wp} (together with the associated template PT_{lib}) such that $PD_{new} = PD$ and $goal_{PR} = goal_{PR_{new}}$ and with $init_{PR_{new}} \vdash w_{PT}$, then PT_{lib} is executable respect to PR_{new} (and therefore

with respect to C_{new}). This makes our templates reusable in a variety of different situations, in which we don't have complete information about the starting state. At this point, the process designer may decide to execute through an external PMS the template PT_{lib} just found, or to refine D_{new} and C_{new} if PT_{lib} does not fit with the designer expectations (for example, if the designer wants a template with a higher degree of concurrency).

Example 6.4.5. Let us suppose that the template shown in Fig. 6.5(a) does not satisfy at all the process designer, since s/he could add one further robot *rb2* to the scenario in order to increase the degree of parallelism in the tasks execution. It follows that a new starting condition $init_{C_3}$ including also contextual information about *rb2* can be defined. The associated initial planning state $init_{PR_3}$, together with the original goal condition $goal_{PR_1}$ and the planning domain PD_1 are first used for verifying if a previously executed template is already stored in the library. The library returns the template PT_1 shown in Fig. 6.5(a), since its weakest preconditions w_{PT_1} are satisfied by $init_{PR_3}$ (i.e., $init_{PR_3} \vdash w_{PT_1}$), and goal condition and planning domain are the same as before. Even if the template in Fig. 6.5(a) is executable with $init_{PR_3}$, the designer may try to search for another plan that (maybe) could exploit the presence of the new robot *rb2*. The planner builds a new plan starting from $init_{PR_3}$, and the associated template PT_2 is shown in Fig. 6.5(b). PT_2 requires the presence of one more robot (i.e., robot *rb2*) and more contextual information for being executed (so its weakest preconditions w_{PT_2} are “richer” than w_{PT_1}), but it provides a higher degree of concurrency in the execution of its tasks. This means that the process designer can choose which template is the best for her/his purposes: one with less concurrency in the tasks enactment but with the fewest participants (cf. Fig. 6.5(a)), or one with more concurrency but requiring more resources for being executed (cf. Fig. 6.5(b)).

Despite the fact that a template is executable “as is”, it can be seen as an “intermediate version” of a completely defined process. In fact, the present POP-based tool cannot be used to synthesize templates involving loops or branching on conditions, and the designer may develop these manually by customizing the template to the specifics of the situation. In the future, one could experiment with more complex AI-planning systems to handle such cases.

6.4.2 Properties

A process template PT guarantees some interesting properties, such as the *correctness* of the template with respect to the information available in the starting state, and the property of *sound concurrency*, meaning that concurrent activities of a template are proven to be effectively independent one from another (i.e., they can not affect the same data).

Definition 23. A process template PT is **correct** respect to a domain theory D and a case C iff it is executable in C .

Theorem 5. Given a solution plan P , a process template PT synthesized for P using our approach is **correct** for any process case C that satisfies the weakest preconditions w_{PT} inferred from P .

The proof of Theorem 1 is straightforward. By definition, a sound planner generates a *consistent* plan [140] that leads from an initial state to a goal. Since we represent the domain theory D and the case C respectively as PDDL planning domain and problem, the planner synthesizes a plan (i.e., a process template) that is correct respect to Definition 23. Note that we could also show that the process template guarantees the achievement of the goal.

A second property we can prove on templates is *sound concurrency*. Despite the fact that in a process designed through the rules imposed by data patterns [111] and workflow patterns [136] the concurrent execution of two or more tasks should guarantee the consistency of data accessed by the concurrent tasks, in practice this is often not true. In fact, in complex environments there isn't a clear correlation between a change in the context and corresponding process changes, making difficult to design by hand a process template where concurrent tasks are also independent. On the contrary, all concurrent tasks of a template synthesized with our approach are proven to be *independent* one from another.

Theorem 6. *Given a process template PT synthesized with our approach, all concurrent tasks are **independent**.*

Proof. By contradiction, let us suppose that a process template PT has two concurrent tasks t_1 and t_2 such that $t_1 \not\parallel t_2$. Hence, t_1 (or t_2) has some effect affecting the precondition of t_2 (or of t_1). This means that $t_1 \triangleright t_2$ or $t_2 \triangleright t_1$. Since PT has been synthesized as result of a POP planner, this dependency between t_1 and t_2 would be represented with a causal link $t_1 \xrightarrow{e} t_2$ (or $t_2 \xrightarrow{e} t_1$), where e is an effect of task t_1 and a precondition for task t_2 (or vice-versa). This causal link requires an ordering between t_1 and t_2 , meaning they need to be executed (and represented in the process template) in sequence. But this means that t_1 and t_2 are not concurrent tasks, by contradicting the original hypothesis. \square

6.5 Translation Algorithms

This section is focussed primarily on presenting technical details concerning the translation algorithms introduced in Section 6.4. Specifically, in the following, we analyze:

- the module PC2PR used for translating a domain theory D and a process case C into a planning domain PD and a planning problem PR ;
- the module POP2PT used for converting a partially ordered plan P into a process template PT .

6.5.1 Representing Domain Theories and Process Cases in PDDL

To obtain a process template that handles a Case C , a corresponding PDDL planning problem definition PR has to be specified. This can be done by mapping $init_C$ to $init_{PR}$ and $goal_C$ to $goal_{PR}$. The planning domain PD is built starting from the definition of ground atomic terms and data types as shown in Section 6.3, and by making explicit the *actions* associated with each annotated task $t \in T$, together with their pre-conditions, effects and input parameters. Basically, the planning

domain describes how predicates and functions change after an action's execution, and specifies the contextual properties constraining the execution of tasks stored in the tasks repository.

Our framework provides a software module named PC2PR in charge of performing such a translation, which makes use of PDDL version 2.1 (cf. [41]). In the following, we discuss how the domain theory D can be translated into a PDDL file representing the planning domain PD:

- the *name* and the *domain* of a data type correspond to an *object type* in the planning domain;
- boolean terms have a straightforward representation as *relational predicates* (templates for logical facts) in the planning domain;
- integer terms correspond to PDDL *numeric fluents*, and are used for modeling non-boolean resources (e.g., the battery charge level of a robot) in the planning domain;
- functional terms do not have a direct representation in PDDL v2.1, but may be represented as relational predicates. Since a functional term is a function $f : Object^n \rightarrow Object$ that maps tuples of objects with domain types D^n to objects with co-domain type U , it may be encoded in the planning domain as a relational predicate P of type (D^n, U) ;
- a given task, together with the associated pre-conditions, effects and input parameters, is translated into a PDDL *action schema*. An action schema describes how the relational predicates and/or numeric fluents change after the action's execution. For example, given the following XML specification of the task $Go \in T$ (with respect to the language provided in Section 6.3):

```
<task>
<name>Go</name>
<parameters>
  <arg>prt - Participant</arg>
  <arg>from - Location</arg>
  <arg>to - Location</arg>
</parameters>
<precondition>at[prt] == from AND
  provides[prt,movement] == true
</precondition>
<effect>at[prt] = to</effect>
</task>
```

the module PC2PR produces the following PDDL representation:

```
(:action go
:parameters (?prt - participant ?from - location
  ?to - location)
:precondition (and ((at ?prt ?from)
  (provides ?prt movement)))
:effect (and (forall (?loc - location) (not (at ?prt ?loc)))
  (at ?prt ?to))
```

This task can be executed only if the actor denoted by *prt* is not currently located in the target location *to* (and is located in his/her starting location *from*) and is able to move into the area. The desired effect turns the value of the predicate $at(prt, to)$ to *true* and $at(prt, from)$ to *false*, meaning the actor moved to the new location.

The planning problem PR is built by translating $init_C$ into $init_{PR}$ and $goal_C$ into $goal_{PR}$:

- for each data type defined in the planning domain, all the possible object instances of that particular data type are explicitly instantiated as *constant symbols* in the initial state of the planning problem (e.g., the fact that $act1$, $act2$, $act3$, $act4$, $rb1$ and $rb2$ are *Participants*, and that $loc00$, ..., $loc33$ are *Locations*);
- a representation of the *initial state* of the planning problem is needed. Basically, the initial state of the planning problem $init_{PR}$ is composed of a conjunction of relational predicates (representing functional and boolean terms in $init_C$) and the initial value of each numeric fluent (e.g., the value of the battery charge level for each robot), corresponding to the values of integer terms in $init_C$;
- the *goal condition* of the planning problem is a logical expression over facts, which partially specifies the state to be reached after the execution of the process template. Again, it is a condition represented as a conjunction of relational predicates and numeric fluent atoms representing the specific boolean, functional and integer terms we want to make true through the correct execution of the process template (as defined in $goal_C$).

6.5.2 Translating a Partially Ordered Plan P into a Process Template PT

Once the plan P has been synthesized, it needs to be translated in a process template PT. As explained in Section 6.2 , a **solution plan** is a three-tuple $P = (A, O, CL)$, where A is the set of actions appearing in the plan, O and CL are respectively the set of ordering constraints and of causal links over A . Since the set of actions A composing the plan and the set of ordering constraints O over A require to be explicitly expressed as nodes and transitions of the template's control flow (as well as their intrinsic ordering), we have implemented a module named POP2PT² that takes as input a solution plan P and converts it into a process template PT.

We provide two algorithms - named respectively "*Find_{PREC/NEXT}*" (cf. Algorithm 1) and "*Build_{PT}*" (cf. Algorithm 2) - to be executed sequentially for automatically computing process template PT. For each planning action $a_i \in A$, Algorithm 1 is in charge of detecting which actions "directly" precede and follow a_i in the plan. This information will be crucial for instantiating the transitions between the flow objects of the process template. However, this knowledge is not directly available in O . In fact, an ordering constraint $a \prec b$ between two actions $a \in A$ and $b \in A$ indicates that a must be executed sometime before action b , but not necessarily immediately before. Algorithm 1, for each action $a_i \in A$, builds two sets containing the actions that immediately precede a_i (the set $PREC(a_i)$) and immediately follow a_i (the set $NEXT(a_i)$).

Definition 24. *Given an ordering constraint $(a \prec b) \in O$ between two actions $a \in A$ and $b \in A$, we say that **a directly precedes b** and **b directly follows a** iff no further action $c \in A$ exists such that $a \prec c$ and $c \prec b$.*

²The software implementing POP2PT is available at http://www.dis.uniroma1.it/~marrella/documents/phd_thesis/templates/POP2PT.zip




Figure 6.6. Overview of the working of the Algorithm $Find_{PREC/NEXT}$.

Basically, for each ordering constraint $o_k = (a_i \prec a_j)^3 \in O$, the algorithm $Find_{PREC/NEXT}$ works as follows:

- if there does not exist any planning action $a_h \neq a_j$ that precedes a_i - i.e., such that $(a_h \prec a_i) \in O$ - then the only predecessor of a_i is the dummy start action a_0 . Therefore, a_0 is added to the set of predecessors of a_i (i.e., $PREC(a_i) = PREC(a_i) \cup \{a_0\}$) and a_i is added to the set of successors of a_0 (i.e., $NEXT(a_0) = NEXT(a_0) \cup \{a_i\}$).
- if there does not exist any planning action $a_h \neq a_i$ that follows a_j - i.e., such that $(a_j \prec a_h) \in O$ - then the only successor of a_j is the dummy end action a_∞ . Therefore, a_j is added to the set of predecessors of a_∞ (i.e., $PREC(a_\infty) = PREC(a_\infty) \cup \{a_j\}$) and a_∞ is added to the set of successors of a_j (i.e., $NEXT(a_j) = NEXT(a_j) \cup \{a_\infty\}$).
- if there does not exist any planning action $a_h \neq a_i$ that precedes a_j - i.e., such that $(a_h \prec a_j) \in O$ - then a_i directly precedes a_j (and a_j directly follows a_i), meaning that $PREC(a_j) = PREC(a_j) \cup \{a_i\}$ and $NEXT(a_i) = NEXT(a_i) \cup \{a_j\}$.
- if there exists a planning action $a_h \neq a_i$ that precedes a_j - i.e., such that $(a_h \prec a_j) \in O$ - but there does not exist any finite sequence of actions a_1, a_2, \dots, a_n such that $(a_i \prec \dots \prec a_1 \prec a_2 \prec \dots \prec a_n \prec \dots \prec a_h)$, then a_i directly precedes a_j (and a_j directly follows a_i), meaning that $PREC(a_j) = PREC(a_j) \cup \{a_i\}$ and $NEXT(a_i) = NEXT(a_i) \cup \{a_j\}$.

Starting from the two sets just computed, the algorithm $Build_{PT}$ is in charge of building the process template PT, by instantiating the tasks, gateways, events and transitions between them. For every action $a \in A$, a corresponding task instance $t_a \in T$ (basically, a task instance is an occurrence of a specific task $t \in T$) is generated (cf. the function *taskify* in Algorithm 2). Transitions between tasks depend on the number of predecessors/successors contained in $PREC(a)/NEXT(a)$. In particular, if an action $b \in A$ is such that $b \in PREC(a)$ or $b \in NEXT(a)$, it is clear that there must be some kind of transition between t_a and t_b in PT:

- if b is the only successor of a , and a is the only predecessor of b (cf. Fig. 6.6(a)), then $(t_a \rightarrow t_b) \in L_T$ (cf. Fig. 6.6(b));
- if a is the only predecessor of b , but b is not the only successor of a (cf. Fig. 6.6(c)), then a parallel split $w_{PS} \in W_{PS}$ is needed between t_a and t_b . Hence $(t_a \rightarrow w_{PS}) \in L_T$ and $(w_{PS} \rightarrow t_b) \in L_{W_{PS}}$ (cf. Fig. 6.6(d));

³Before the execution of the algorithm, all the ordering constraints involving the dummy start action a_0 and the dummy end action a_∞ are removed from O .

- if b is the only successor of a , but a is not the only predecessor of b (cf. Fig. 6.6(e)), then a parallel join $w_{PJ} \in W_{PJ}$ is needed between t_a and t_b . Hence $(t_a \rightarrow w_{PJ}) \in L_T$ and $(w_{PJ} \rightarrow t_b) \in L_{W_{PJ}}$ (cf. Fig. 6.6(f));
- if t_a is not the only predecessor of t_b and t_b is not the only successor of t_a (cf. Fig. 6.6(g)), then a parallel split $w_{PS} \in W_{PS}$ and a parallel join $w_{PJ} \in W_{PJ}$ are needed between t_a and t_b . Hence $(t_a \rightarrow w_{PS}) \in L_T$, $(w_{PS} \rightarrow w_{PJ}) \in L_{W_{PS}}$ and $(w_{PJ} \rightarrow t_b) \in L_{W_{PJ}}$ (cf. Fig. 6.6(h)).

Finally, if an action $a \in A$ has no predecessors/successors (i.e., the set $PREC(a)/NEXT(a)$ is empty), this means that t_a must be connected with the start event \circ and end event \odot .

6.6 Experiments

To show the feasibility of the approach, we ran some experiments and measured the time required for synthesizing a partially ordered plan for some variants of our running example described in Section 6.1. We ran our tests using POPF2 [19], which is a temporal planner that handles PDDL 2.1 [41] and preserves the benefits of partial-order plan construction in terms of producing makespan-efficient, flexible plans. Search in POPF2 is based around the idea of expanding a partial-order plan in a forward direction; steps added to the plan are ordered after a subset of those in the partial plan, rather than after every step in the plan so far. We conducted the experiments on a Intel U7300 CPU 1.30GHz Dual Core, 4GB RAM machine.

The experimental setup was run on variants of the test case shown in our running example. We represented 7 planning actions in PD (corresponding to 7 different emergency management tasks stored in the tasks repository \mathbb{T}), annotated with 7 relational predicates and 6 numeric fluents, in order to make the planner search space sufficiently challenging. Then, we defined 18 different planning problems of varying complexity by manipulating the number of facts in the goal. As well, we examined how irrelevant domain knowledge affects the performance of the planner. Starting from a planning problem PR with an initial state $init_{PR}$ completely specified and with a goal condition $goal_{PR}$ expressed as the conjunction of n facts, we manipulated the specification of the initial state $init_{PR}$ to reduce the number of known facts. In our experiments, the number of facts in goal condition ranges from 1 single fact to a conjunction of 6 logical facts (that make the contextual problem harder). As shown in Table 1, for a given goal condition composed of n facts, our purpose was to measure the computation time needed for finding a sub-optimal solution for problems specified with starting states with a decreasing amount of knowledge. The column labeled as “Knowledge in $init_{PR}$ ” makes explicit which information is removed from the initial state of the planning problem with a given goal condition $goal_{PR}$. For example, if we consider our running scenario from Section 6.1, whose goal condition is composed of 3 facts and characterized by a complete specification of the starting state, the time needed for finding a solution plan is of 0.13 seconds. After removing from the initial state all the information concerning the actor act_3 , the time required for computing the plan decreases to 0.11 seconds. In general, for a given goal condition, removing “irrelevant information” from the initial state reduces the search space and the computation required for synthesizing the plan. It is also

Algorithm 1: $Find_{PREC/NEXT}$ - Find actions predecessors and successors

Data:

- A : the set of actions appearing in the final plan P , including dummy actions a_0 and a_∞ .
- O : the set containing the ordering constraints returned by the planner, in the form $o_k = (a_i \prec a_j)$, with $(a_i, a_j) \in A$. Ordering constraints that involve a_0 and a_∞ are removed.
- $NEXT(a_i)$: a set containing the list of successors of the i -th action $a_i \in A$.
- $PREC(a_i)$: a set containing the list of predecessors of the i -th action $a_i \in A$.
- k : an integer number, used as counter for the ordering constraints.
- $length(S)$: returns the size of a set S .
- $add(a, S)$: inserts an action a into a set S .

Result: for each $a_i \in A$, it returns $NEXT(a_i)$ and $PREC(a_i)$

Init :

```

    k = 0
    for each  $a_i \in A$  do
        NEXT( $a_i$ ) =  $\emptyset$ 
        PREC( $a_i$ ) =  $\emptyset$ 
begin
    while  $k < length(O)$  do
        take the  $k$ -th ordering constraint,  $o_k = (a_i \prec a_j)$  from  $O$ 
        start scanning the  $O$  set
        if  $\nexists (a_h \in A) : a_h \neq a_j \wedge (a_h \prec a_i \in O)$  then
            add( $a_i, NEXT(a_0)$ );
            add( $a_0, PREC(a_i)$ );
        if  $\nexists (a_h \in A) : a_h \neq a_i \wedge (a_j \prec a_h \in O)$  then
            add( $a_\infty, NEXT(a_j)$ );
            add( $a_j, PREC(a_\infty)$ );
        if  $\nexists (a_h \in A) : a_h \neq a_i \wedge (a_h \prec a_j \in O)$  OR
            $\exists (a_h \in A) : a_h \neq a_i \wedge (a_h \prec a_j \in O)$  AND  $\nexists$  any finite sequence of
           actions  $a_1, a_2, \dots, a_n$  such that
           ( $a_i \prec \dots \prec a_1 \prec a_2 \prec \dots \prec a_n \prec \dots \prec a_h$ ) then
            add( $a_j, NEXT(a_i)$ );
            add( $a_i, PREC(a_j)$ );
        else
            do nothing, because it means that  $a_i \prec a_h \prec a_j$ . The ordering
            constraint  $a_i \prec a_h$  will be considered in a future iteration of the
            algorithm.
    k++

```

Algorithm 2: *Build*_{PT} - Build Process Template**Data:**

- A : the set of actions appearing in the final plan, including a_0 and a_∞ .
- $L = L_T \cup L_E \cup L_{W_{PS}} \cup L_{W_{PJ}}$ is a finite set of transitions connecting events, task instances and gateways. Initially it is empty.
- $NEXT(a_i)/PREC(a_i)$: a set with the list of successors/predecessors of a_i .
- $length(S)$: returns the size of a set S .
- $insert(x, S)$: given an element x and a set S , if $x \notin S$ inserts x into S .
- $taskify(a)$: given a planning action $a \in A$, it generates a corresponding task instance $t_a \in T$.

```

begin
  for each  $a_i \in A$  do
     $t_{a_i} = taskify(a_i)$ 
    if  $length(PREC(a_i)) > 1$  then
      for each  $a_q \in PREC(a_i)$  do
         $t_{a_q} = taskify(a_q)$ 
        if  $length(NEXT(a_q)) > 1$  then
          |  $insert(\{w\_split_q \rightarrow w\_join_i\}, L_{W_{PS}})$ 
        else
          |  $insert(\{t_{a_q} \rightarrow w\_join_i\}, L_T)$ 
       $insert(\{w\_join_i \rightarrow t_{a_i}\}, L_{W_{PJ}})$ 
    else
      if  $length(NEXT(a_i)) > 1$  (let  $a_q$  be the only predecessor of  $a_i$ ) then
        if  $length(PREC(a_q)) > 0$  then
          |  $insert(\{w\_split_q \rightarrow t_{a_i}\}, L_{W_{PS}})$ 
        else
          |  $insert(\{\circ \rightarrow w\_split_q\}, L_E)$ 
          |  $insert(\{w\_split_q \rightarrow t_{a_i}\}, L_{W_{PS}})$ 
        else
          if  $length(PREC(a_q)) > 0$  then
            |  $insert(\{t_{a_q} \rightarrow t_{a_i}\}, L_T)$ 
          else
            |  $insert(\{\circ \rightarrow t_{a_i}\}, L_E)$ 
      if  $length(NEXT(a_i)) > 1$  then
        for each  $a_j \in NEXT(a_i)$  do
           $t_{a_j} = taskify(a_j)$ 
          if  $length(PREC(a_j)) > 1$  then
            |  $insert(\{w\_split_q \rightarrow w\_join_i\}, L_{W_{PS}})$ 
          else
            |  $insert(\{w\_split_i \rightarrow t_{a_j}\}, L_{W_{PS}})$ 
           $insert(\{t_{a_i} \rightarrow w\_split_i\}, L_T)$ 
      else
        if  $length(PREC(a_j)) > 1$  (let  $a_j$  be the only successor of  $a_i$ ) then
          if  $length(NEXT(a_j)) > 0$  then
            |  $insert(\{t_{a_i} \rightarrow w\_join_j\}, L_T)$ 
          else
            |  $insert(\{t_{a_i} \rightarrow w\_join_j\}, L_T)$ 
            |  $insert(\{w\_join_j \rightarrow \circ\}, L_{W_{PJ}})$ 
        else
          if  $length(NEXT(a_j)) > 0$  then
            |  $insert(\{t_{a_i} \rightarrow t_{a_j}\}, L_T)$ 
          else
            |  $insert(\{t_{a_i} \rightarrow \circ\}, L_T)$ 

```

Table 6.1. Time performances of POPF2.

Facts in $goal_{PR}$	Knowledge in $init_{PR}$	Time for a sub-opt. sol.
1	complete state	0.17
	No information about act1	0.15
	No information about act1 and act3	0.12
2	complete state	0.12
	No information about act3	0.10
	No information about act3 and rb1	0.08
3	complete state	0.13
	No information about act3	0.11
	No information about act3 and rb2	0.09
4	complete state	0.21
	No information about act3	0.20
	No information about act3 and rb1	0.10
5	complete state	0.17
	No information about act3	0.16
	No information about act3 and rb1	0.10
6	complete state	1.56
	No information about act3	1.19
	No information about act3 and act1	1.13

interesting to note that a sub-optimal solution includes more actions than those strictly required for fulfilling the goal condition. In particular, when the number of facts in a goal condition increases, the quality of the sub-optimal solution produced typically decreases. Based on our experiments, the approach seems quite feasible for medium-sized dynamic processes as used in practice.

6.7 Related Work

Process modeling is the first and most important step in the BPM lifecycle [142], which intends to provide a high-level specification of a business process that is independent from implementation and serves as a basis for process analysis, automation, and verification. The task of defining a model is often performed with the aid of tools that provide a graphical representation, but without any automatic generation of the process model. However in recent years, numerous AI planning-based approaches have been devised for the latter, and the closest to our approach are [43, 99, 40]. [43] presents the basic idea behind the use of planning techniques for generating a process schema, but no implementation seems to be provided, and the direct use of the PDDL language for annotating tasks and specifying the domain theory requires a deep understanding of AI planning technology.

In [99], the authors exploit the IPSS planner [101] for modeling processes in SHAMASH [3], a knowledge-based system that uses a rule-based approach. To automate the process model generation, they first translate the semantic representation of SHAMASH into the IPSS language. Then, IPSS produces a parallel plan of activities that is finally translated back into SHAMASH and is presented graphically to the user. This work proposes the scheduling of parallel activities (that implicitly handle time and resource constraints), meta-modeling that deals with planning explicitly, and suggests that learning could be used for process optimization. However, the emphasis here is on supporting processes for which one has complete knowledge. This assumption does not always hold for dynamic processes, where some contextual information may not be available at the time of process model

synthesis. The work of [40] proposes a new life cycle for workflow management based on the continuous interplay between learning and planning. The approach is based on the use of machine learning algorithms for inferring pre-conditions and effects of activities, and generate a partially-ordered execution plan (i.e., a process model with concurrent branches) that complies to these rules. An interesting result concerns the possibility of producing process models even though the activities may not be accurately described. In such cases, the authors use a best-effort planner that is always able to create a plan, even though the plan may be incorrect. By refining the preconditions and effects of planning actions, the planner will be able to produce several candidate plans, and after a finite number of refinements, the best candidate plan (i.e., the one with the lowest number of unsatisfied preconditions) may be chosen and translated into a process model. Unfortunately, this approach has not been very successful, as the best plan typically moves further and further away from the correct solution.

6.8 Conclusion

In this chapter, we developed a technique based on partial-order planning algorithms and declarative specifications of process tasks for synthesizing a library of concurrent process templates to be enacted in partially specified contextual scenarios. Characteristic of these processes is the role of contextual data acting as a driver for process modeling. We are currently working on a complete implementation and thorough validation of the whole approach, including the formalization of metrics for evaluating process templates' quality. A future direction for this work is to generate hierarchical process templates, with high-level templates achieving more general goals that can invoke simpler templates to achieve some of their subgoals. We also plan to address expressiveness limitations, such as handling preferences, representing negative preconditions (negative literals in goals and preconditions are not supported by most POP planners, including POPF2) and tasks with context-dependent effects.

Chapter 7

Recovering Dynamic Processes in YAWL

The need to automatically adapt processes in response to exceptions, events and contextual changes has emerged as a leading issue, both in dynamic and pervasive scenarios and in processes with a varying degree of structuring (cf. Section 1.2). In order to demonstrate the general validity of the approach for automatic adaptation proposed in Chapter 3, a further direction of the author's research activity was to integrate such an approach into YAWL [124], that is among the most well-known PMSs coming from academia. To this end, our contribution has been twofold:

- In [80] we contextualize and demonstrate our adaptation approach in the service-oriented environment provided by the YAWL system. We leverage the *Flexibility as a Service* approach [129] to obtain a complete integration between YAWL and SmartPM by designing and implementing a so-called YAWL Custom Service [133], named the SmartPM Service. At design time, the process designer is thus able to associate atomic tasks in the YAWL specification with the SmartPM Service. Atomic tasks defined in YAWL may be decomposed into complex sub-processes defined for the SmartPM environment. At runtime the service is then able to check-out a work-item and execute the corresponding subprocess. When a subprocess is executed in SmartPM, the automatic adaptivity features are exploited in order to react to exogenous events and to progress the subprocess. Upon completion of a subprocess, control is passed back to the YAWL environment along with any data that was produced.
- In [81] we rely on the modeling capabilities provided by the YAWL language and on the process enactment and exception handling capabilities provided by the YAWL environment, and we propose significant extensions at both the modeling and architectural level. Specifically, we introduce and define PLANLETS, as self-contained YAWL specifications where process tasks are annotated at design-time with pre-conditions, desired effects and post-conditions. In the presence of an exception, this approach allows delegating to an external planner the automatic run-time synthesis of a suitable recovery procedure by contextually selecting the compensation tasks from a specific repository linked to the PLANLET under execution.




Figure 7.1. The YAWL process defined for a train derailment scenario (a), in which the composite task “Manage Emergency in the Area” is a dynamic process (b).

For describing the two approaches, we have considered again a motivating example drawn from an emergency management setting (cf. Section 7.1). It clearly represents an application scenario for which it is unrealistic to assume that all exceptional situations, as well as required exception handlers, can be anticipated at design-time and thus be incorporated a priori into the pre-specified process model.

The rest of the chapter is organized as follows. In Section 7.2 we briefly focus on the exception handling approach implemented in the YAWL system, as presented in [124]. In Section 7.3 we show how the YAWL environment (and its imperative modeling approach) can be complemented with our SmartPM execution environment that exploits a declarative modeling approach to deal with environmental changes and exceptions in process executions. In Section 7.4 we provide an in-depth discussion and concrete design and implementation proposal of how the YAWL architecture can be extended to support PLANLETS and to integrate planning capabilities. Finally, Section 7.5 discusses limitation and future developments of the approach.

7.1 Running Example

As an application scenario, we consider again an emergency management process defined for train derailments. The corresponding YAWL process to be executed is shown in Figure 7.1(a). The process starts when the railway traffic control center receives an accident notification from the train driver and collects some information about the derailment, including the train ID code, the GPS location and the number of coaches and passengers. Then it could be required to cut off the power in the area and to interrupt the railway traffic near the derailment scene. In parallel, after having collected additional information about the train (e.g., security equipment) and about emergency services available in the area, an emergency response team can be sent to the derailment scene.

Collected information is used for defining and configuring at run-time an incident response plan, defined by a contextually and dynamically selected set of activities to be executed on the field by first responders. Such activities are abstracted into the




Figure 7.2. The YAWL architecture.

composite task¹ “Manage Emergency in the Area” (cf. Fig. 7.1(b)). The subnet is composed by three parallel branches with tasks that instruct first responders to act for evacuating people from train coaches, to take pictures and to assess the gravity of the accident.

Note that the YAWL sub-process defined in Figure 7.1(b) corresponds to the same process used as case study in Section 1.3. For the sake of simplicity, we also consider the same contextual information and scenario shown in Figure 1.2(b) and described in Section 1.3.

7.2 The YAWL Architecture

The YAWL system, like the majority of classical PMSs, allows to define stable and well-understood processes and offer adequate support as long as the processes are structured and do not require much flexibility. The exception handling capabilities provided by YAWL were designed and implemented starting from the conceptual framework for workflow exception handling presented in [113]. In order to understand how exceptions are detected and handled in YAWL we refer to the architecture² in Figure 7.2.

At design-time, the process designer identifies the possible events that may result in deviations from the expected process execution at run-time, and defines both the detection and exception handling strategies. To date, YAWL³ is able to deal with eight different types of exceptions, i.e., case-level and workitem-level pre/post-execution constraints violations, case-level and workitem-level externally triggered exceptions, timeouts and unavailability of resources.

¹A composite task is a container for another YAWL sub-net, with its own set of elements.

²The picture refers to the architecture defined in [124].

³We refer to the final release of YAWL 2.1.

For each exception that can be anticipated, it is possible to define an exception handling process, named *exlet*, which includes a number of exception handling primitives (for removing, suspending, continuing, completing, failing and restarting a workitem/case) and one or more compensatory processes in the form of *worklets* (i.e., self-contained YAWL specifications executed as a replacement for a workitem or as compensatory processes). Exlets are linked to specifications by defining specific *rules* (through the *Rules Editor* graphical tool), in the shape of Ripple Down Rules specified as *if condition then conclusion*, where the *condition* defines the exception triggering condition and the *conclusion* defines the exlet.

At run-time, exceptions are detected and managed by the *Exception Service*, a sub-service of the *Worklet Service*. The *Exception Service* is notified by the engine via *Interface X* of exception triggering events (which include timeouts, resource unavailabilities and notifications fired when a case/workitem begins/ends in order to check pre/post-execution constraints) that may result in an exception⁴. For each event notification, the service determines whether an exception has occurred and, if so, it executes the corresponding exlet. Exception handling primitives are directly executed invoking the corresponding methods (for removing, suspending, etc. a workitem/case) provided by the engine-side of *Interface X*. If the exlet includes a compensation worklet, the *Exception Service* first retrieves it from the repository, then loads it into the engine via *Interface A* and finally starts it via *Interface B*. The worklet is then executed by the engine as a new separate case, possibly in parallel with the parent case if it was not suspended by the exlet.

7.3 Making YAWL and SmartPM interoperate

The service-oriented approach that characterizes the architecture of YAWL makes the system easily extendable and provides direct support for implementing the *Flexibility as a Service* approach [129]. In YAWL the engine manages running cases, but is not directly responsible for task executions. Resources, entities and systems able to execute tasks are abstracted as *services* that interact with the YAWL engine via a set of interfaces. Specifically, the interaction between the engine and the Custom Services mainly occurs through Interface B (cf. Figure 7.2).

At design-time each atomic task in a YAWL specification can be associated with a so-called Custom Service [133] (e.g., the Default Worklist Handler/Resource Service, the Worklet Service, the Declare Service, etc.) that at run-time is responsible for task execution according to its internal logic. To this end, we designed and implemented a Custom Service, named **SmartPM Service**, that allows the YAWL execution environment to delegate the execution of sub-processes and activities to the **SmartPM** execution environment which is able to automatically adapt a process to deal with emerging changes and exceptions. As shown in Figure 7.3, the **SmartPM Service** allows the interaction between the two environments.

The **SmartPM Service** enables the decomposition of YAWL tasks into processes to be executed by **SmartPM**. At design-time, the process designer is thus able to associate atomic tasks in the YAWL specification with the **SmartPM Service**. The

⁴Externally triggered exceptions are instead notified by the environment, specifically by a client involved in process execution.




Figure 7.3. The integration between YAWL and SmartPM.

service requires as input variable a process to be performed by SmartPM, defined according to the formalism introduced in Chapter 3. If the SmartPM process is already available at design-time, it can be directly associated with the input variable required by the service. However, a process to be delegated to SmartPM can also be built starting from an available template, which is configured and finalized by exploiting data produced as output by other tasks in the main YAWL process. In this case, the executable SmartPM process has to be produced as output by a YAWL task that precedes the task associated with the SmartPM Service. The domain-dependent configuration of a SmartPM process can be done either manually by a process designer or automatically by a dedicated service.

In the example shown in Figure 7.1(a), we have an high-level process for managing railway accidents, which has been modeled in and is currently managed by YAWL. One of the tasks, specifically the one where a team of operators has to act on the field for evacuating people from train coaches (cf. the complex activity “Manage Emergency in the Area” in Figure 7.1(a)), represents a step in the process model referring to a complex dynamic process (cf. Figure 7.1(b)). As thoroughly described in Section 1.3, we assume that operators are equipped with mobile devices, and aspects like context-awareness and managing frequent exogenous events (e.g., bad connections between devices and operators) is crucial. For this reason, the high dynamism of the operating environment requires that such activities are executed by the SmartPM environment, which provides automatic adaptation features. Therefore, every time the complex activity “Manage Emergency in the Area” gets enabled during process execution, the enactment of its corresponding sub-process model is delegated to the SmartPM environment through the SmartPM service.

As any other Custom Service, the SmartPM service implements the service-side of Interface B and is thus able to receive notifications from the YAWL engine when a new work-item is created and delegated to the service for execution. Invoking the specific methods provided by the engine-side of Interface B, the SmartPM Service is then able to check-out a work-item (i.e., notify the engine that the work-item is going to be executed) and execute the corresponding SmartPM input process. Once the execution of the subprocess has been completed, the service can check-in the work-item (i.e., notify the engine of execution completion), again via Interface B, along with the corresponding output. Specifically, the SmartPM Service produces as output, among other possible values, a boolean one that indicates whether the input subprocess was successfully completed or not. Control is then passed back to the YAWL environment, which can continue to carry out the main process.

A screen-cast of a demonstration that shows the interaction between the two environments is available at <http://www.dis.uniroma1.it/~marrella/public/DemoBPM2011.zip>. The application presented in the demo is currently a proof of concept and aimed at demonstrating the feasibility and validity of the approach. In the demo application some tasks managed by YAWL and the devices of the operators are simulated, whereas the integration among the two systems and the communication between SmartPM and autonomous devices is real. The next step towards the definition and implementation of a full-fledged solution has been achieved in [81] and is described in the following section, where we discuss how concretely the adaptation approach shown in Chapter 3 can be built on top of YAWL.

7.4 The Planlets Approach

In this section we propose a solution that builds on top of YAWL [124], and consists in annotating at design-time a YAWL specification with additional information which allows process instances to be automatically recovered. In particular, we assume the tasks of a YAWL process specification to be annotated with *pre-conditions*, *desired effects* and *post-conditions*. Failures arise either when associated pre-conditions for a task are not satisfied at the time the task is to be started, or when post-conditions do not hold after the execution of the task. Effects represent the changes that a successful task execution imposes on the *state* of the world, reflecting the current value of the contextual properties that constraint the process under execution. Hence, the process designer just states *what* conditions have to be satisfied, without having to anticipate *how* these can be fulfilled. In order to formalize the concept, we introduce the definition of PLANLET:

Definition 25 (Planlet). *Let YN be a YAWL net, T be the tasks defined in YN , and V be the set of variables defined in YN . Let $Expr(V)$ be the set of expressions over the variables in V . A PLANLET is a tuple $(YN, Pre, Post, Eff)$ where*

- *$Pre : T \rightarrow Expr(V)$ returns an expression representing the pre-conditions of tasks in T ;*
- *$Post : T \rightarrow Expr(V)$ returns an expression representing the post-conditions of tasks in T ;*
- *$Eff : T \rightarrow Expr(V)$ returns an expression representing the effects of tasks T .*

The role of pre/post-conditions and effects for a YAWL task is twofold:

1. pre-conditions and post-conditions enable run-time process execution monitoring and exception detection: they are checked respectively before and after task executions, and the violation of a pre-condition or post-condition results in an exception to be handled;
2. along with the input/output parameters consumed/produced by the task, pre-conditions and effects provide a complete specification of the task: this allows the task to be represented as an action in a planning domain description and therefore used for solving a planning problem built to handle an exception.

At design-time, the annotated tasks are stored in a repository linked to the PLANLET specification, which may contain also other annotated tasks deriving from previous executions on the same contextual domain. At run-time, while instances of the YAWL specification are carried on, tasks become enabled. Every time a task $t \in T$ becomes enabled, expression $Pre(t)$ is evaluated; similarly, upon the completion of t , expression $Post(t)$ is evaluated. If an evaluation returns false upon enablement or completion of a task, the system is in an *invalid state* and, hence, the YAWL specification instance needs to be adapted to come back into the “right track”. In order to do that, the case execution is suspended, and a recovery procedure is automatically synthesized.

To provide more details, let us assume that the current PLANLET is $\delta_0 = (\delta_1; \delta_2)$ in which δ_1 is the part of the PLANLET already executed and δ_2 is the part of the PLANLET which remains to be executed when an exception is identified. The adapted PLANLET is $\delta'_0 = (\delta_1; \delta_h; \delta_2)$. However, whenever a PLANLET needs to be adapted, every running task is interrupted, since the “repair” sequence of tasks $\delta_h = [t_1, \dots, t_n]$ is placed before them. Thus, active branches can only resume their execution after the repair sequence has been executed. This last requirement is fundamental to avoid the risk of introducing data inconsistencies during a repair.

The automatic synthesis of the recovery procedure δ_h is enacted on-the-fly by an external planner. A planner solves the problem to find a sequence of actions that move a system state from the initial one to a target goal, using a predefined set of admissible actions. Each action is associated the set of pre-conditions $Pre(t)$ in order for that step to be chosen, as well as the effects $Eff(t)$ obtained as result of the action’s execution.

Along with defining the set of admissible actions, it is also crucial to define how the state is represented, since pre-conditions and effects of actions are given in term of the chosen state representation. The actions’ set and the state definition are often referred to as *planning domain*. The standard representation language of planners to define actions and state is the Planning Domain Definition Language (PDDL) [46]. In the context of adaptation of instances of YAWL specifications, each task specification is associated with a different action in the planning domain; the task’s pre-conditions and effects are translated in PDDL and associated to the corresponding action. In addition to the so-created planning domain, when an exception arises, the invalid state and the pre-condition (or post-condition) violated is given in input to a planner, which can try to build a plan. If the plan exists, the planner is eventually going to return it. In this case, the plan is converted into a sequence of YAWL tasks which are assigned to qualifying participants. When




Figure 7.4. The YAWL architecture extended with the *Planning Service*.

the converted plan is carried out, the original suspended process is restored for execution.

Example 7.4.1. Let us consider the example introduced in Section 7.1. The composite activity “Manage Emergency in the Area” may be modeled as a self-contained PLANLET specification (cf. Fig. 7.1(b)), linked to a repository containing a set of emergency management (annotated) tasks, that range from the simple activity of taking pictures to the more complex extinguishment of a fire. An explicit representation of contextual information (the connection of each actor to the network, the map of the area, the battery level of each robot etc.) is needed for preserving the correct PLANLET execution. Suppose now that the task $go(loc00, loc33)$ is assigned to actor $act1$, which reaches instead the location $loc03$ (cf. Fig. 1.3(b)). This means that the actor $act1$ is not more connected to the network and s/he is in a position different than the desired one, resulting in a *post-condition failure*. The planner builds a planning problem by taking as initial state the invalid state of the PLANLET, and as goal a state where all actors/robots are inter-connected to the network and $act1$ is in the desired location $loc33$. The recovery plan is automatically synthesized by contextually selecting tasks from the repository linked to the PLANLET. Suppose, for example, that the two robots $rb1$ and $rb2$ have an empty battery. In such a case, the planner devises on-the-fly a possible solution, composed by a sequence of 5 tasks [$chargeBattery(act4, rb1)$, $move(rb1, loc00, loc03)$, $go(act1, loc03, loc33)$, $chargeBattery(act4, rb2)$, $move(rb2, loc00, loc33)$] that change the state of the world as shown in Fig. 1.4(d). Note that the recovery plan is synthesized by taking care of the skills of process participants, and their availability for task assignment and execution. Hence, each task composing the plan is already associated to the participant that will execute it.

7.4.1 Incorporating Planlets into YAWL

The architectural extension and integration we designed takes advantage of YAWL’s exception detection capabilities and leverages the flexibility of the exlet-based




Figure 7.5. *Planning Service* activation hierarchy for exception handling.

handling techniques. From an architectural perspective, as shown in Fig. 7.4, planning capabilities are provided by a *Planning Service* that implements the planning logic and algorithm. In order to define the role of the *Planning Service* and clarify how it interacts with existing YAWL architectural components and services, we follow the process and exception handling life-cycle, from process design, enactment and monitoring to exception detection, handling and (possibly) resolution. At design time, the process designer builds one or more *PLANLET Repositories* (or modifies the existing ones), by inserting/deleting annotated tasks and by (possibly) modifying the contextual domain linked to each repository. Tasks involved in a *PLANLET* specification are selected from a specific *PLANLET* repository, since they are thought to be enacted in a specific contextual domain. Before to execute a *PLANLET*, the process designer instantiates the initial values for the properties of the contextual domain. As shown in Section 7.4.2, tasks pre- and post-conditions are automatically translated in YAWL pre- and post-constraints. In order to delegate the exception handling to the *Planning Service*, we introduce the possibility of mapping a *compensation activity* to the *Planning Service*. By defining this mapping instead of explicitly selecting a compensation worklet, the process designer configures the *Exception Service* so that the generation of the compensation worklet is delegated to the *Planning Service*. Fig. 7.5 shows an excerpt of the rule file defined for detecting and handling a workitem-level pre-execution (or post-execution) constraint violation. Lines 1-4 define the exception triggering condition (a pre- or a post-condition failure), while lines 5-12 define the exception handling exlet (which consists in suspending the current case, performing some compensation activities and then resuming the suspended case). In our extended version, the mapping of a compensation task to the *Planning Service* is identified by a `<target>` element containing the `PlanningService` value (line 9), in order to enact planning capabilities.

Example 7.4.2. If we consider our running example, the compensation plan devised in Fig. 7.5 corresponds to the one needed for re-establishing the network connection between actors/robots and for instructing actor `act1` to move in the desired location.

Planning Service Activation. When the *Exception Service* activates the *Planning*

Service, it provides as input all case data associated with the running case, along with the detected violation over pre- and post-conditions. Based on this information, and on the specifications of available tasks, stored in the repository linked to the PLANLET under execution, the *Synchronization* component of the *Planning Service* is able to build the planning domain and to define a planning problem, and submit them to the *Planner* module in charge of synthesizing a recovery plan. If the *Planner* is able to successfully synthesize a compensation plan, it stores it as an executable specification (i.e., a worklet) in the *Worklets Repository* and notifies the *Exception Service*. The *Exception Service* is then able to enact the execution of the compensation worklet as if it was manually selected at design time, by loading the specification into the engine and launching it as a separate case. When the execution completes, output data produced by the worklet are mapped back to the parent case and subsequent actions in the exlet are executed. Following the exlet defined in Fig. 7.5, as the compensation worklet synthesized by the *Planner* is supposed to recover from the constraint violation, the suspended case can then be resumed and executed. If no valid plan can be found by the *Planner*, a notification alert is sent to an administrator, who is in charge of handling the unsolved exception, e.g., manually building a compensation process or just canceling the process case.

7.4.2 Annotating YAWL Specifications in Planlets

A main step of our approach in YAWL consists in enriching the process model with a specification of process tasks, in terms of pre-conditions, desired effects and post-conditions, and with an explicit representation of the contextual domain needed for the correct process enactment.

In YAWL, each atomic task t can be linked to a decomposition. Decompositions can have a number of input and output *parameters*, each identified by a *name* and characterized by a *type* dictating valid values it may store, and define the so-called YAWL Service that will be responsible for task execution. As process data are represented through net-level variables, inbound and outbound mappings define how data is transferred from net variables to task variables and vice-versa. We propose to extend task specifications at the decomposition level, with the possibility of defining pre-conditions, post-conditions and effects as logical formulae and expressions over task parameters.

Defining and representing finite domain types. The definition of a PLANLET requires the specification of the data types that characterize the information manipulated by process instances and define the domains over which predicates and functions are interpreted. In order to have a compact and finite representation of a process state, given by the values assumed by process variables at a given point in the execution, all data types must correspond to *finite* domains over which variables of that type can range; this requirement is imposed by the planning-based approach we propose. Examples of such domains are finite integer intervals or sets of strings, and other enumerated domains. As YAWL applies strong data typing and all data types are defined using XML Schemas, this can be easily achieved by defining data types as XML Schemas and using restrictions (e.g., via the `enumeration` constraint) to limit the content of an XML element to a set of acceptable values.

In our example, we need to define data types for representing actors, robots⁵ and locations in the area. The following listing provides the (simplified) definition of the data type $Loc = \{loc00, loc10, \dots, loc33\}$.

```
<xs:simpleType name="Loc">
  <xs:restriction base="xs:string">
    <xs:enumeration value="loc00"/>
    ...
    <xs:enumeration value="loc33"/>
  </xs:restriction>
</xs:simpleType>
```

Under this representation, we consider possible values as constant symbols that univocally identify objects in the domain of interest; different variables of the same type having the same value denote the same object.

Defining and representing predicates and functions. The annotation of tasks with preconditions, postconditions and effects requires the definition of the set of predicates, numeric functions and object functions to be used when extending task definitions. Predicates and functions have to be completely specified at design time. Predicates can be used to express properties of domain objects and relations over objects. A predicate consists of a predicate *symbol* P and a set of *typed parameters* or *arguments*⁶. Argument types (taken from the set of data types previously defined) represent the finite domains over which predicates are interpreted. A generic predicate declaration is of the form:

$$P(arg_1 : type_1, arg_2 : type_2, \dots, arg_n : type_n)$$

In our example, we may need predicates for expressing the presence of a fire in a location or whether a location is covered by the network signal provided by the main antenna, or relations, such as the adjacency between locations, i.e.,

$$Fire(loc : Loc) \quad Covered(loc : Loc) \quad Adjacent(loc1 : Loc, loc2 : Loc)$$

In addition to basic predicates, we allow the designer to define *derived* predicates. They are declared as basic predicates, with the additional specification of a well-formed formula φ that determines the truth value for the predicate :

$$P(arg_1 : type_1, \dots, arg_n : type_n) \{ \varphi \}$$

In our domain, we may need to express that an actor is connected to the network if s/he is in a covered location or if s/he is in a location adjacent to a location where a robot is located (and is thus connected through the robot); assuming we have

⁵Although emergency operators and robots can be considered as *resources* or *services* able to execute tasks and can be represented in the organizational model provided by YAWL, we also need to explicitly represent them in the process because we need to define predicates and functions over these domains.

⁶Predicates with no arguments, i.e., with arity 0, are allowed and can be considered as propositions; they are directly represented as boolean variables.

defined the data types $Robot = \{rb1, rb2\}$ and $Actor = \{act1, act2, act3, act4\}$, we have:

$$Connected(act : Actor) \{ EXISTS(l1 : Loc, l2 : Loc, rbt : Robot) (\\ (at(act) = l1) AND (Covered(l1) OR (atRobot(rbt) = l2 AND Adjacent(l1, l2)))) \}$$

Numeric and Object Functions. Numeric and object functions allow to represent and handle numeric values and domain objects as functions of other objects. Function declarations are of the form:

$$f(arg_1 : type_1, arg_2 : type_2, \dots, arg_n : type_n) : DataType$$

and consist of a function *symbol* f , a set of *typed parameters*⁷, and a *return type*. Numeric functions have as return type an integer or a real number, whereas object functions have a return type taken from the set of data types defined in the net specification. The arguments of functions range over *finite* domains, and for object functions the same requirement holds for result types (i.e., object functions must have finite co-domains). In our example, we need to keep track of the battery level of the robots. This can be represented through the numeric function

$$batteryLevel(robot : Robot) : Integer \quad \text{and} \quad batteryStep : Integer$$

Similarly, we can represent the position of actors and robots by defining the following functional predicates that map actors and robots to their location:

$$at(actor : Actor) : Loc \quad \text{and} \quad atRobot(robot : Robot) : Loc$$

State variables representation. The use of predicates and functions requires that at run-time we represent the corresponding logical interpretations, as state variables that hold (a) the truth value of the defined predicates over domain objects, and (b) the values of the defined functions with respect to different argument assignments. The interpretations are used to evaluate pre- and post-conditions, and are modified as a result of task executions. As a consequence of the declaration of a predicate or function, two new data types are automatically generated and added to the XML data types definitions for the net:

T1. a complex data type that is able to represent the name of the predicate or function and

- for predicates, all argument assignments for which the predicate holds⁸ (i.e., the current interpretation P^I for the predicate);
- for functions, all argument assignments for which the function is defined, along with the corresponding value⁹ (i.e., the current interpretation f^I for the function);

⁷Numeric functions with no arguments are allowed, and can be considered as state variables rather than constants, as their value may change during process executions; they are represented as integer or float/double variables.

⁸Basically, a set containing all object tuples for which the predicate is true.

⁹Basically, a map where object tuples are mapped to objects.

- T2. a complex data type that is able to represent a predicate or function instance, in terms of the name of the predicate or function, the set of arguments and their assignment, and the truth value or numeric/object value of the predicate or function with respect to the specific assignment; different parameters of this type can be defined for process tasks, to be used for representing the effects that they can have on the predicate or function interpretation.

For each predicate and function, a single net-level state variable of type T1 is defined and it can be initialized so as to contain all values for the objects for which the predicate is true or the function is defined in the initial state. Derived predicates are not explicitly represented through net-level state variables, as their interpretation can be always derived from the corresponding formula, and they can not appear in task effects (but task effects can act on the basic predicates and functions that appear in the formula, thus indirectly modifying the truth value for the derived predicate).

Initial interpretation. The initial interpretation over which a process instance is executed is given by an assignment of values to the state variables that represent truth values for predicates and values functions. The initial interpretation is thus defined by a set of grounded predicates (initial facts) and initialization values for functions.

The following listings provide an example of a possible initialization for the variables of type T1 for the *Adjacent* predicate and the *batteryLevel* function, representing the grounded predicates (facts) *Adjacent(loc00, loc01)*, *Adjacent(loc00, loc10)* and so on, and the assignments *batteryLevel(rb1) = 2* and *batteryLevel(rb2) = 4*.

```

<PredicateAdjacent>
  <name>Adjacent</name>
  <interpretation>
    <loc1>loc00</loc1>
    <loc2>loc01</loc2>
  </interpretation>
  <interpretation>
    <loc1>loc00</loc1>
    <loc2>loc10</loc2>
  </interpretation>
  ...
</PredicateAdjacent>

<FunctionBatteryLevel>
  <name>BatteryLevel</name>
  <interpretation>
    <robot>rb1</robot>
    <value>2</value>
  </interpretation>
  <interpretation>
    <robot>rb2</robot>
    <value>4</value>
  </interpretation>
</FunctionBatteryLevel>

```

Pre-conditions, post-conditions and effects. Pre-conditions, post-conditions and effects are defined at design time as logical annotations associated with tasks in a PLANLET. Basically, we assume a first-order predicate logic with numeric and object functions, with the restriction that free variables are not allowed and thus all variable symbols must be task parameter names or occur in the scope of a quantifier. The language is clearly inspired from PDDL, although we prefer an infix notation for the operators.

Pre and postconditions can be defined as an atomic formula as well as the conjunction (AND) or disjunction (OR) of formulae, and formulae can be negated

(NOT) and existentially or universally quantified (EXISTS and FORALL).

Atomic formulae are (basic or derived) predicates defined over argument terms, where a term can be a task parameter or a constant; the equality ($==$) or negated equality ($!=$) predicates between terms are supported. In addition, it is possible to define conditions as expressions that make use of relational binary comparison operators ($<$, $>$, $=$, \leq , \geq) and involve numeric expression. Numeric expressions are constructed, using arithmetic operators ($+$, $-$, $*$, $/$), from primitive numeric expressions, which include integers and real numbers, and numeric functions defined over argument terms. Numeric expressions are not allowed to appear as terms, i.e., as arguments to predicates or values of task parameters (otherwise, predicates would be defined over arguments with *infinite* ranges). Object functions are defined over argument terms, and can be used as terms only in equality and inequality atomic formulae.

Task effects basically define the changes that a successful task execution imposes on the current state of the world, as represented in the process execution context. A task effect can be defined as an atomic formula as well as the conjunction of formulae, and atomic formulae can be negated. As for preconditions, atomic formulae are predicates defined over argument terms (task parameters or constants), but derived predicates and equality are not allowed. In addition to boolean predicates, task effects may include assignment expressions to update the values of numeric functions. A numeric effect consists of an assignment operator, the numeric function to be updated and a numeric expression. Assignment operators include (i) direct assignment ($=$), to assign to a numeric function a numerical value defined by a number or by a numeric expression; (ii) relative assignments, which can be used to increase ($+=$) or decrease ($-=$) the value of a numeric function (additive assignment effects), as well as to scale-up ($*=$) or scale-down ($/=$) the value of a numeric function (scaling assignment effects). Similarly, it is possible to update the value of object functions with an assignment operator, restricted to direct assignment ($=$), to assign to an object function a typed object identified by a term.

Marking up task specifications. Task pre-conditions, post-conditions and effects are represented in task specifications via the `<precondition>`, `<effect>` and `<postcondition>` markup elements. In our example, consider the task labeled as `go`, which require that an actor moves from a location to another in the area. It defines two input parameters *from* and *to* of type *Loc*, representing the starting and arrival locations, and an input parameter *actor* of type *Actor* representing the emergency operator that executes the task. An instance of this task can be executed only if s/he is currently at the starting location and is connected to the network. As a consequence (effect) of task execution, the actor moves from the starting to the arrival location, but we need, as post-condition, to verify whether the arrival location has been reached and the actor is still connected to the network. We can thus define the following annotations:

```
<precondition>at(actor) == from AND Connected(actor)</precondition>
<effect>at(actor) = to</effect>
<postcondition>at(actor) == to AND Connected(actor)</postcondition>
```

The designer can distinguish between: (i) *direct effects*, i.e., effects that always take place after an execution, and therefore the corresponding changes on the state

variables are automatically performed when the task completes (e.g., if an effect of the form $BatteryLevel(robot) += 5$ is marked as automatic, after task execution the value for $BatteryLevel(robot)$ is directly increased by 5); and (ii) *supposed* effects, i.e., effects that define changes that are assumed to be performed only when the task is considered as an action in a planning domain. Supposed effects can be interpreted as the effects that a task is supposed to have, but the actual produced changes are defined at run-time as a result of the concrete execution, such as the actual truth value of a predicate or the actual value for a direct assignment.

Example 7.4.3. In our example, $at(actor) = to$ is a supposed effect, as the actual value for $at(actor)$ is produced as a task output and may be different from the desired one (i.e., the value of the to variable prescribed in the effect). If the designer needs to verify whether a task execution has produced the intended effect, s/he has to define a corresponding post-condition (i.e., the $at(actor) == to$).

Direct effects can be directly represented by generating an outbound mapping with an XQuery expression that adds/removes a tuple to/from the state variable representing predicate's interpretation (for positive/negative predicates), or updates the value for a tuple in the state variable representing function's interpretation (for assignment effects). In supposed effects, the actual values are produced by workitem executions, and all predicates and functions that appear in the effect expression have to be represented as task variables, so as to allow to specify (according to task's execution logic) the truth value for predicates or the value for functions. To this end, we represent each predicate and function that appears in the supposed effects as task parameters of type T2, where the predicate/function name is given and fixed, the values for the argument variables (i.e., the grounding) are defined by the inbound mappings for task parameters and the predicate/function actual value will be defined as a result of task execution. For these variables, outbound mappings are then generated, including XQuery expressions to update net-level state variables as for direct effects.

Example 7.4.4. Fig. 7.6 shows an example of how a variable can be used to represent an effect and how the actual value for $at(act1)$ can be produced as output; we show that the output value for $at(act1)$ is produced by a sensor (i.e., a GPS device) supporting the worklist handler. The produced value, in the example 'loc33', is then used to update the net variable representing the at interpretation to reflect that $at(act1) \mapsto loc33$.

State model and exceptions In a PLANLET, a process state S is given by the token marking m_S (as defined in [124] for YAWL nets) and the logical interpretation \mathcal{I}_S that assigns truth values to predicates and values to functions. The initial state over which a process instance is executed is given by the initial marking and an assignment of values to the state variables that represent the initial interpretation for predicates and functions. When a task t becomes enabled in a state S (as determined by m_S), its execution can start only if the task precondition formula φ_{pre} is true in \mathcal{I}_S , i.e., $\mathcal{I}_S \models \varphi_{pre}$. A task execution changes the interpretation according to actual task effects (which for a successful execution are given by the corresponding effects expression $expr_{eff}$) and leads to a new state S' where $m_{S'}$ is the produced marking and $\mathcal{I}_{S'}$ is the new interpretation. A completed task is considered as successfully

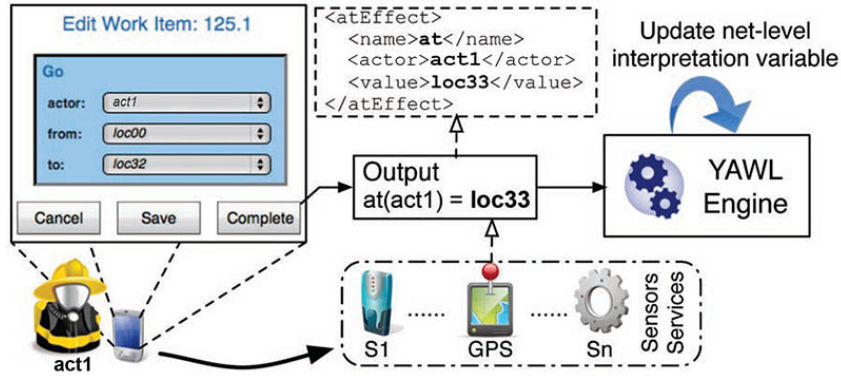


Figure 7.6. A task effect represented as a variable assignment.

executed if its postcondition formula φ_{post} is true in $\mathcal{I}_{S'}$, i.e., $\mathcal{I}_{S'} \models \varphi_{post}$. At run time all task executions are thus preceded and followed by the verification of whether $\mathcal{I} \models \varphi_{pre}$ and $\mathcal{I} \models \varphi_{post}$ ¹⁰. In this model, an exception occurs in a given state with an interpretation \mathcal{I} if a task is enabled but $\mathcal{I} \not\models \varphi_{pre}$ or if a task has completed but $\mathcal{I} \not\models \varphi_{post}$.

From pre-/post-conditions to pre-/post-execution constraints As part of its exception handling mechanism, YAWL supports the definition of workitem-level pre- and post-execution constraints, as rules with conditions that (i) are checked when the workitem becomes enabled and when it is completed, and (ii) if violated, they trigger an exception and the execution of an exception handling process (i.e., a YAWL exlet). Conditions are defined over case variables as strings of operands and arithmetic, comparison and logical operators; conditional expressions may also take the form of boolean XQuery expressions [124]. In our approach, we leverage on this built-in feature and map the evaluation of pre- and post-conditions to the evaluation of pre and post-execution constraints, by automatically translating φ_{pre} and φ_{post} formulae for each task into YAWL conditional expressions. While arithmetic, comparison and logical operators in our annotation language directly map to the operators supported by YAWL, predicates and functions can be resolved by appropriate XQuery expressions, according to a standard evaluation algorithm for boolean queries¹¹; basically:

- all variables in the formula that correspond to task parameters (for which an inbound or outbound mapping from or to a net variable is defined as an XQuery expression) assume the values as defined by the mappings
- each function, which becomes defined over a tuple of ground terms, is replaced by an XQuery expression that extract from function's interpretation state variable the value corresponding to the tuple (i.e., $f(c_1, \dots, c_k)$ is replaced by the value $f^{\mathcal{I}}(c_1, \dots, c_k)$);
- each predicate, which becomes defined over a tuple of ground terms, is replaced by a boolean XQuery expressions that verify whether the tuple is contained in

¹⁰As φ_{pre} and φ_{post} are closed formulae, their truth values can be considered as the answers to the corresponding boolean queries, given the interpretation \mathcal{I} .

¹¹we recall that no free variables are allowed and all formulae are closed

the predicate interpretation state variable ($\mathcal{I} \models P(c_1, \dots, c_k)$ if $(c_1, \dots, c_n) \in P^{\mathcal{I}}$)

- existentially and universally quantified formulae are replaced by boolean XQuery expressions that verify whether there exist an assignment for the quantified typed variable that makes the formula true, or whether the formula is true for all possible assignments (among the finite possible values defined in the data type specification)
- derived predicates in preconditions are replaced by the corresponding formulae, which are evaluated as per the previous items

Representing Planlet Annotations in PDDL. In order to exploit our planning-based recovery mechanism, every task/annotation/property associated to a PLANLET needs to be translated in PDDL. A PDDL definition consists of two parts: the domain and the problem definition. The planning domain is built starting by the definition of basic/derived predicates, object/numeric functions and data types as shown in the previous sections, and by making explicit the *actions* associated to each annotated task stored in the repository linked to the PLANLET under execution, together with the associated pre-conditions, effects and input parameters. Basically, the planning domain describes how predicates and functions may vary after an action execution, and reflects the contextual properties constraining the execution of tasks stored in a specific PLANLET repository. Our annotation syntax allows to represent planning domains and problems with the complexity of those describable in PDDL version 2.2 [37], that is characterized for enabling the representation of realistic planning domains.

In the following, we discuss how our annotations are translated into a PDDL file representing the planning domain:

- the *name* and the *domain* of a data type corresponds to an *object type* in the planning domain;
- basic and derived predicates have a straightforward representation as *relational predicates* (templates for logical facts) and *derived predicates* (to model the dependency of given facts from other facts) in the planning domain;
- numeric functions correspond to PDDL *numeric fluents*, and are used for modeling non-boolean resources (e.g., the battery level of a robot) in the planning domain;
- object functions do not have a direct representation in PDDLv2.2, but may be replaced as relational predicates. Since an object function $f : Object^n \rightarrow Object$ map tuples of objects with domain types D^n to objects with co-domain type U , it may be coded in the planning domain as a relational predicate P of type (D^n, U) ;
- a given YAWL task, together with the associated pre-conditions and effects and input parameters, is translated in a PDDL *action schema*. An action schema describes how the relational predicates and/or numeric fluents may vary after the action execution. In the following, it is shown the PDDL representation of the task *go*:

```
(:action go
:parameters (?x - actor ?from - loc ?to - loc)
```

```


```

:precondition (and (not (at ?x ?to)) (at ?x ?from) (connected ?x))
:effect (and (not (at ?x ?from)) (at ?x ?to))

```


```

This task can be executed only if the actor denoted with x is not currently located in the target location to (and, consequently, is located in his/her starting location $from$) and is connected to the network. The desired effect turns the value of the predicate $at(x, to)$ to TRUE and $at(x, from)$ to FALSE, meaning the actor moved in a new location.

When an exception arises, on a same planning domain a new planning problem is built at run-time, through the description of an initial state (that corresponds to the invalid state s of the process) and the description of the desired goal (a safe state s' , derived from the violated pre- or post-condition).

- for each data type defined in the planning domain, all the possible object instances of that particular data type are explicitly instantiated as *constant symbols* in the planning problem (e.g., the fact that $act1, act2, act3, act4$ are *Actors*, $rb1$ and $rb2$ are *Robots*, $loc00, \dots, loc33$ are *Locations*);
- a representation of the *initial state* of the planning environment is needed. Basically, the initial state of the planning problem corresponds to an invalid state (i.e., a state that needs to be fixed after a pre- or post-condition violation during the process execution). It is composed by a conjunction of relational predicates, derived predicates (e.g., the information about which actors/robots are currently connected to the network) and by the current value of each numeric fluent (e.g., the value of the battery level for each robot);
- the *goal state* of the planning problem is a logical expression over facts. In our approach, the goal state is built in order to reflect a safe state to be reached after the execution of a recovery procedure. Suppose that t is the task whose pre-conditions $Pre(t)$ (or post-conditions $Post(t)$) are not verified. The safe state s' corresponding to the goal state of the planning problem is generated starting from the invalid state s , by substituting the wrong facts that led to the exception with the content of the pre-conditions (or post-conditions) violated.

Example 7.4.5. In the case of our running example, actor $act1$ reaches a wrong location $loc03$ after the execution of the task go , and s/he is no longer connected to the network, resulting in a post-condition failure (cf. Figure 1.3(b)). Therefore, whereas the initial state of the planning problem generated for dealing with such an exception corresponds exactly to the invalid state of the process (hence, with the value of the predicate $Connected(act1)=FALSE$ and $at(act1)=loc03$), the goal state is composed by all those predicates/functions not affected by the task go (that must remain unchanged after the recovery procedure), and by the post-conditions just violated, i.e., $Connected(act1)=TRUE$ and $at(act1)=loc33$, which is the desired location.

7.5 Conclusion

In this chapter, we have customized our general approach for automatic adaptation of dynamic processes presented in Chapter 3 by discussing its deployment on top of

existing PMSs.

Specifically, we have first shown how the YAWL environment can be complemented with the SmartPM execution environment by leveraging the “Flexibility as a Service” approach [129]. A screen-cast of a demonstration that shows the interaction between the two environments is available at <http://www.dis.uniroma1.it/~marrella/public/DemoBPM2011.zip>.

Then, we have presented a concrete design and implementation proposal of how the YAWL architecture can be extended to integrate planning capabilities. For this aim, we have proposed the approach of PLANLETS, self-contained YAWL specifications with recovery features, based on modeling of pre- and post-conditions of tasks and the use of planning techniques. The feasibility of the PLANLET approach has been investigated by performing some testing to learn the time amount needed for synthesizing a recovery plan for different adaptation problems. We made our tests by using the LPG-td planner [45, 44], and the obtained results can be compared with the ones shown in Section 5.1.

Currently, the implementation of the PLANLETS approach is ongoing, in collaboration with researchers from the Queensland University of Technology (QUT, Australia).

Chapter 8

Conclusion

The research activity outlined in this thesis has been devoted to define a general approach, a concrete architecture and a prototype PMS for automatic adaptation of dynamic processes, on the basis of a declarative specification of process tasks and relying on well-established planning techniques. Our purpose was to demonstrate that the combination of procedural and imperative models with declarative elements, along with the exploitation of techniques from the field of artificial intelligence (AI) such as planning algorithms and tools, can increase the ability of existing PMSs of supporting dynamic processes.

To this end, we developed a prototype PMS named *SmartPM*, which is specifically tailored for supporting collaborative work of process participants during pervasive scenarios. The adaptation mechanism deployed on *SmartPM* is based on execution monitoring for detecting failures at run-time, which does not require the definition of the adaptation strategy in the process itself (as most of the current approaches do), and on automatic planning techniques for the synthesis of the recovery procedure.

In order to exploit the automatic adaptation features provided by *SmartPM*, some extra modeling effort is required at design time. In fact, our approach requires that processes are defined at best partly textually (for the formalization of the domain theory associated to the dynamic process) and partly graphically (for the definition of the process control flow). However, we think the overhead is compensated at run-time by the automation of adaptation procedures. While, in general, such modelling effort may seem significant, in practice it is comparable to the effort needed to encode the adaptation logic using alternative methodologies like happens, for example, in rule-based approaches.

The assumptions of classical planning (determinism in the action effects, the closed-world assumption for the instantiation of the initial situation, etc.) we used for modeling dynamic processes has a twofold consequence. On the one hand, we can exploit the good performance of classical planners (e.g., LPG-td [45], POPF2 [19]) to solve real-world problems with a realistic complexity. In fact, test results reported in Section 5.1 show that the time overhead introduced by a planner for synthesizing recovery procedures with variable length for adaptation problems of growing complexity (in terms of size and parameters of the corresponding planning problems) is in the order of tens of seconds, which makes the approach feasible for medium-sized dynamic processes used in practice. On the other hand,

classical planning imposes some restrictions for addressing more expressive problems, including incomplete information, preferences and multiple task effects. Future works will include an extension of our approach dealing with the above aspects and by considering also classical business scenarios, with the purpose to maintain the planning process very responsive.

A second future work concerns to integrate the approach for building process templates (cf. Chapter 6) in the SmartPM system. In fact, when considering specific classes of emergency management processes (e.g., earthquakes, floods, building collapses, etc.), it is possible to identify recurring and predefined activities that have to be performed according to general guidelines and emergency action plans. In order to balance between the need to create ad-hoc process specifications and to exploit guidelines and recurrent activity patterns, it is possible to take advantage of our approach for generating process templates. Processes can be designed on-demand from provided templates, which are tailored and configured to create an appropriate process specification. This kind of approach would allow to speed up process definition activities, as well as to facilitate the reuse of process fragments.

Finally, even if our approach is able to adapt a process instance at run-time, it does not allow to evolve the original process model on the basis of exceptions captured. Therefore, a third future work concerns to avoid to consider all deviations from the process as errors, but as a natural and valuable part of the work activity, which provides the opportunity for learning and thus evolving the process model for future instantiations.

Appendix A

The Full Code of the Example

This appendix is focussed primarily for describing the full code of the example related to our case study. We show, in sequence, the SmartML specification linked to the dynamic process depicted in Section 1.3, the IndiGolog program to be executed by the PMS and the PDDL Planning Domain and Problem built for dealing with the exception detected in Fig. 1.3. We underline that the IndiGolog program reflecting the approach shown in Chapter 3 is written with the ProLog language.

The SmartML specification of the domain theory

```
TYPES = {Participant, Capability, Location_type, Status_type, Boolean_type, Integer_type}
```

```
Participant = {act1,act2,act3,act4,rb1,rb2}
Capability = {movement,hatchet,camera,gprs,extinguisher,battery,digger,powerpack}
Location_type = {loc00,loc10,loc20,loc30,loc01,loc11,loc02,loc03,loc13,loc23,loc31,
                loc32,loc33}
Status_type = {ok,fire,debris}
Boolean_type = {true,false}
Integer_type = {0..30}
```

```
PRE-DEFINED TERMS = {provides, requires}
```

```
provides[prt:Participant] = (bool:Boolean_type)
requires[prt:Participant] = (bool:Boolean_type)
```

```
NON RELEVANT TERMS = {atRobot, batteryLevel, photoTaken, neigh, covered, generalBattery,
batteryRecharging, moveStep, debrisStep}
```

```
atRobot[prt:Participant] = (loc:Location_type)
batteryLevel[prt:Participant] = (int:Integer_type)
photoTaken[loc:Location_type] = (int:Integer_type)
neigh[loc1:Location_type,loc2:Location_type] = (bool:Boolean_type)
covered[loc:Location_type] = (bool:Boolean_type)
generalBattery[] = (int:Integer_type)
batteryRecharging[] = (int:Integer_type)
moveStep[] = (int:Integer_type)
debrisStep[] = (int:Integer_type)
```

```
RELEVANT TERMS = {evacuated, at, status}
```

```

at[prt:Participant] = (loc:Location_type)
evacuated[loc:Location_type] = (bool:Boolean_type)
status[loc:Location_type] = (st>Status_type)

COMPLEX TERMS = {isConnected, isRobotConnected}

isConnected(prt:Participant) {
  EXISTS(l1:Location_type, l2:Location_type, rbt:Participant).((at(act)=l1) AND
  (covered(l1) OR (atRobot(rbt)=l2 AND neigh(l1,l2)=true AND isRobotConnected(rbt))))}.

isRobotConnected(rbt:Participant) {
  EXISTS(l1:Location_type, l2:Location_type).((atRobot(rbt)=l1) AND
  (covered(l1) OR
  (neigh(l1,l2) AND Covered(l2)) OR
  (EXISTS(l3:Location_type, l4:Location_type,
  l5:Location_type, rbt2:Participant).atRobot(rbt2)=l3 AND
  ((neigh(l1,l3) AND covered(l3)) OR
  (neigh(l1,l5) AND neigh(l3,l5) AND covered(l3)) OR
  (neigh(l3,l4) AND covered(l4) AND neigh(l1,l3)) OR
  (neigh(l3,l4) AND covered(l4) AND neigh(l1,l5) AND neigh(l3,l5)))))).}

EXOGENOUS EVENTS = {photoLost, fireRisk, rockSlide}

<ex-events>
  <ex-event>
    <name>rockSlide</name>
    <parameters>
      <arg>loc - Location_type</arg>
    </parameters>
    <effects>
      <automatic>status[loc] = debris</automatic>
    </effects>
  </ex-event>
  <ex-event>
    <name>fireRisk</name>
    <parameters>
      <arg>loc - Location_type</arg>
    </parameters>
    <effects>
      <automatic>status[loc] = fire</automatic>
    </effects>
  </ex-event>
  <ex-event>
    <name>photoLost</name>
    <parameters>
      <arg>loc - Location_type</arg>
    </parameters>
    <effects>
      <automatic>photoTaken[loc] = false</automatic>
    </effects>
  </ex-event>
</ex-events>

TASKS REPOSITORY = {go, move, takephoto, evacuate, updatestatus,
  extinguishfire, chargebattery}

```



```

<tasks>
  <task>
    <name>go</name>
    <parameters>
      <arg>from - Location_type</arg>
      <arg>to - Location_type</arg>
    </parameters>
    <precondition>at[PRT] == from AND isConnected[PRT] == true</precondition>
    <effects>
      <supposed>at[PRT] = to</supposed>
    </effects>
  </task>
  <task>
    <name>move</name>
    <parameters>
      <arg>from - Location_type</arg>
      <arg>to - Location_type</arg>
    </parameters>
    <precondition>atRobot[PRT] == from AND batteryLevel[PRT] >= moveStep[] AND
      isRobotConnected[PRT] == true
    </precondition>
    <effects>
      <supposed>atRobot[PRT] = to</supposed>
      <automatic>batteryLevel[PRT] -= moveStep[]</automatic>
    </effects>
  </task>
  <task>
    <name>takephoto</name>
    <parameters>
      <arg>loc - Location_type</arg>
    </parameters>
    <precondition>at[PRT] == loc AND isConnected[PRT] == true</precondition>
    <effects>
      <supposed>photoTaken[loc] = true</supposed>
    </effects>
  </task>
  <task>
    <name>evacuate</name>
    <parameters>
      <arg>loc - Location_type</arg>
    </parameters>
    <precondition>at[PRT] == loc AND isConnected[PRT] == true AND
      status[loc] == ok AND evacuated[loc] == false
    </precondition>
    <effects>
      <supposed>evacuated[loc] = true</supposed>
    </effects>
  </task>
  <task>
    <name>updatestatus</name>
    <parameters>
      <arg>loc - Location_type</arg>
    </parameters>
    <precondition>at[PRT] == loc AND isConnected[PRT] == true AND
      status[loc] == ok
    </precondition>
    <effects>

```

```

        <supposed>status[loc] = ok</supposed>
    </effects>
</task>
<task>
    <name>removedebris</name>
    <parameters>
        <arg>loc - Location_type</arg>
    </parameters>
    <precondition>atRobot[PRT] == loc AND isRobotConnected[PRT] == true AND
        status[loc] == debris AND batteryLevel[PRT] >= debrisStep[]
    </precondition>
    <effects>
        <supposed>status[loc] = ok</supposed>
        <automatic>batteryLevel[PRT] -= debrisStep[]</automatic>
    </effects>
</task>
<task>
    <name>extinguishfire</name>
    <parameters>
        <arg>loc - Location_type</arg>
    </parameters>
    <precondition>at[PRT] == loc AND isConnected[PRT] == true AND
        status[loc] == fire
    </precondition>
    <effects>
        <supposed>status[loc] = ok</supposed>
    </effects>
</task>
<task>
    <name>chargebattery</name>
    <parameters>
        <arg>rb - Participant</arg>
    </parameters>
    <precondition>provides[rb,battery] == true AND at[PRT] == atRobot[rb] AND
        isConnected[PRT] == true AND generalBattery[] >= batteryRecharging[]
    </precondition>
    <effects>
        <automatic>batteryLevel[rb] += batteryRecharging[]</automatic>
        <automatic>generalBattery[] -= batteryRecharging[]</automatic>
    </effects>
</task>
</tasks>

```

The IndiGolog Program

```

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% FILE: aPMS.pl
% DESCRIPTION : The most recent version of the IndiGolog adaptive Process Management System.
% VERSION : 1.0
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

/* SOME PRE-BUILT COMMANDS AND MACROS*/
:- dynamic controller/1.

cache(_):-fail.

```

```

causes_true(_,_,_) :- false.
causes_false(_,_,_) :- false.

actionNum(X,X).

square(X,Y) :- Y is X * X.
member(ELEM,[HEAD|_]) :- ELEM=HEAD.
member(ELEM,[_|TAIL]) :- member(ELEM,TAIL).
listEqual(L1,L2) :- subset(L1,L2),subset(L2,L1).

/* THE LIST OF AVAILABLE SERVICES */
service(S) :- domain(S,[act1,act2,act3,act4,rb1,rb2]).

/* THE LIST OF CAPABILITIES RELEVANT FOR THE PROCESS OF INTEREST */
capability(B) :- domain(B,[movement,hatchet,camera,gprs,extinguisher,battery,digger,powerpack]).

/* THE CAPABILITIES PROVIDED BY EACH SERVICE*/
provides(act1,movement).
provides(act1,gprs).
provides(act1,extinguisher).
provides(act1,camera).
provides(act2,movement).
provides(act2,gprs).
provides(act2,hatchet).
provides(act3,movement).
provides(act3,gprs).
provides(act3,hatchet).
provides(act4,movement).
provides(act4,powerpack).
provides(act4,gprs).
provides(rb1,battery).
provides(rb1,digger).
provides(rb2,battery).
provides(rb2,digger).

/* PRE-DEFINED DATA TYPES */
boolean_type(Q) :- domain(Q,[true,false]).

integer_type(N) :- domain(N,[0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17,18,19,20,
21,22,23,24,25,26,27,28,29,30]).

/* USER-DEFINED DATA TYPES */
status_type(S) :- domain(S,[ok,fire,debris]).

location_type(L) :- domain(L,[loc00,loc01,loc02,loc03,loc10,loc11,loc13,loc20,loc23,loc30,
loc31,loc32,loc33]).

/* TASKS DEFINED IN THE PROCESS SPECIFICATION */
task(T) :- domain(T,[chargebattery,go,move,evacuate,removedebris,takephoto,
updatestatus,extinguishfire]).

/* THE LIST OF ADMISSIBLE IDENTIFIERS */
task_identifiers([id_1,id_2,id_3,id_4,id_5,id_6,id_7,id_8,id_9,id_10,
id_11,id_12,id_13,id_14,id_15,id_16,id_17,id_18,id_19,id_20,id_21,
id_22,id_23,id_24,id_25,id_26,id_27,id_28,id_29,id_30,id_adapt]).

```

```

id(D) :- domain(D,task_identifiers).

/* THE CAPABILITIES REQUIRED BY EACH TASK*/
requires(go,movement).
requires(evacuate,hatchet).
requires(takephoto,camera).
requires(updatestatus,gprs).
requires(extinguishfire,extinguisher).
requires(move,battery).
requires(removedebris,digger).
requires(chargebattery,powerpack).

/* THE LIST OF ADMISSIBLE WORKITEMS */
workitem(go, ID, [FROM, TO], [TO]) :- id(ID),location_type(FROM),location_type(TO).
workitem(takephoto, ID, [LOC], [true]) :- id(ID),location_type(LOC),boolean_type(true).
workitem(evacuate, ID, [LOC], [true]) :- id(ID),location_type(LOC),boolean_type(true).
workitem(move, ID, [FROM, TO], [TO]) :- id(ID),location_type(FROM),location_type(TO).
workitem(chargebattery, ID, [RB], []) :- id(ID),service(RB).
workitem(removedebris, ID, [LOC], [ST]) :- id(ID),location_type(LOC),status_type(ST).
workitem(extinguishfire, ID, [LOC], [ST]) :- id(ID),location_type(LOC),status_type(ST).
workitem(updatestatus, ID, [LOC], [ST]) :- id(ID),location_type(LOC),status_type(ST).

/* TASKS PRECONDITIONS */

prim_action(assign(SRVC, ID, TASK, _INPUTS, _EXOUTPUTS)) :- service(SRVC),id(ID),task(TASK).

poss(assign(SRVC, [workitem(go, ID, [FROM, TO], [TO])]),
      and(isConnected(SRVC),at(SRVC)=FROM)) :-
      service(SRVC),id(ID),location_type(FROM),location_type(TO).

poss(assign(SRVC, [workitem(takephoto, ID, [LOC], [true])]),
      and(isConnected(SRVC),at(SRVC)=LOC)) :-
      service(SRVC),id(ID),location_type(LOC),boolean_type(true).

poss(assign(SRVC, [workitem(evacuate, ID, [LOC], [true])]),
      and(isConnected(SRVC),and(at(SRVC)=LOC,
      and(evacuated(LOC)=false,status(LOC)=ok)))) :-
      service(SRVC),id(ID),location_type(LOC),boolean_type(true).

poss(assign(SRVC, [workitem(move, ID, [FROM, TO], [TO])]),
      and(isRobotConnected(SRVC),
      and(atRobot(SRVC)=FROM,batteryLevel(SRVC)>=moveStep))) :-
      service(SRVC),id(ID),location_type(FROM),location_type(TO).

poss(assign(SRVC, [workitem(chargebattery, ID, [RB], [])]),
      and(isConnected(SRVC), (and(at(SRVC)=atRobot(RB),
      and(provides(RB,battery),generalBattery>=batteryRecharging)))))) :-
      service(SRVC),id(ID),service(RB).

poss(assign(SRVC, [workitem(removedebris, ID, [LOC], [ST])]),
      and(isRobotConnected(SRVC),and(atRobot(SRVC)=LOC,
      and(status(LOC)=debris,batteryLevel(SRVC)>=debrisStep)))) :-
      service(SRVC),id(ID),location_type(LOC),status_type(ST).

poss(assign(SRVC, [workitem(extinguishfire, ID, [LOC], [ST])]),
      and(isConnected(SRVC),and(at(SRVC)=LOC,status(LOC)=fire))) :-
      service(SRVC),id(ID),location_type(LOC),status_type(ST).

```

```

poss(assign(SRVC,[workitem(updatestatus,ID,[LOC],[ST])]),
    and(isConnected(SRVC),and(at(SRVC)=LOC,status(LOC)=ok))) :-
    service(SRVC),id(ID),location_type(LOC),status_type(ST).

/* BASIC ACTIONS AND ENGINE FLUENTS FOR MANAGING THE TASK LIFE-CYCLE */

rel_fluent(assigned(SRVC,ID,TASK)) :- service(SRVC),id(ID),task(TASK).
causes_val(assign(SRVC,ID,TASK,_INPUTS,_EXOUTPUTS),assigned(SRVC,ID,TASK),true,true).
causes_val(release(SRVC,ID,TASK,_INPUTS,_EXOUTPUTS,_PHOUTPUTS),assigned(SRVC,ID,TASK),false,true).

rel_fluent(reserved(SRVC,ID,TASK)) :- task(TASK), service(SRVC), id(ID).
causes_val(readyToStart(SRVC,ID,TASK),reserved(SRVC,ID,TASK),true,true).
causes_val(finishedTask(SRVC,ID,TASK,_),reserved(SRVC,ID,TASK),false,true).

prim_action(start(SRVC,ID,TASK)) :- service(SRVC), id(ID), task(TASK).
poss(start(SRVC,ID,TASK),reserved(SRVC,ID,TASK)) :- service(SRVC),id(ID),task(TASK).

prim_action(ackCompl(SRVC,ID,TASK)) :- task(TASK),service(SRVC),id(ID).
poss(ackCompl(SRVC,ID,TASK),neg(reserved(SRVC,ID,TASK))).

prim_action(release(SRVC,ID,TASK,_INPUTS,_EXOUTPUTS,_PHOUTPUTS)) :- service(SRVC),id(ID),task(TASK).
poss(release(_SRVC,_ID,_TASK,_INPUTS,_EXOUTPUTS,_PHOUTPUTS), true).

rel_fluent(free(SRVC)) :- service(SRVC).
causes_val(release(SRVC,_ID,_TASK,_INPUTS,_EXOUTPUTS,_PHOUTPUTS),free(SRVC),true,true).
causes_val(assign(SRVC,_ID,_TASK,_INPUTS,_EXOUTPUTS),free(SRVC),false,true).

/* IF THIS FLUENT HOLDS, THE SYSTEM SWITCHES FROM USING THE PLAN-BASED ADAPTATION MECHANISM
TO USE THE BUILT-IN ADAPTATION APPROACH */
rel_fluent(built_in_adaptation).

/*USER-DEFINED STATIC FLUENTS*/
fun_fluent(neigh(LOC1,LOC2)) :- location_type(LOC1),location_type(LOC2).
fun_fluent(covered(LOC)) :- location_type(LOC).

/*FLUENT AT*/
fun_fluent(at(SRVC)) :- service(SRVC).

causes_val(release(SRVC,ID,go,_I,_E,[OUTPUT]),at(SRVC),OUTPUT,
    assigned(SRVC,[workitem(go,ID,[FROM,TO],[TO])])) :-
    service(SRVC),id(ID),location_type(OUTPUT),location_type(FROM),location_type(TO).

fun_fluent(at_exp(SRVC)) :- service(SRVC).

causes_val(release(SRVC,ID,go,_I,_E,[OUTPUT]),at_exp(SRVC),TO,
    and(neg(adapting),assigned(SRVC,[workitem(go,ID,[FROM,TO],[TO])])) :-
    service(SRVC),id(ID),location_type(OUTPUT),location_type(FROM),location_type(TO).

/*FLUENT AT_ROBOT*/
fun_fluent(atRobot(SRVC)) :- service(SRVC).

causes_val(release(SRVC,ID,move,_I,_E,[OUTPUT]),atRobot(SRVC),OUTPUT,
    assigned(SRVC,[workitem(move,ID,[FROM,TO],[TO])])) :-
    service(SRVC),id(ID),location_type(OUTPUT),location_type(FROM),location_type(TO).

```

```

/*FLUENT EVACUATED*/
fun_fluent(evacuated(P)) :- location_type(P).

causes_val(release(SRVC, ID, evacuate, _I, _E, [OUTPUT]), evacuated(LOC), OUTPUT,
  assigned(SRVC, [workitem(evacuate, ID, [LOC], [true])])) :-
  service(SRVC), id(ID), location_type(LOC), boolean_type(true), boolean_type(OUTPUT).

fun_fluent(evacuated_exp(P)) :- location_type(P).

causes_val(release(SRVC, ID, evacuate, _I, _E, [OUTPUT]), evacuated_exp(LOC), true,
  and(neg(adapting), assigned(SRVC, [workitem(evacuate, ID, [LOC], [true])])) :-
  service(SRVC), id(ID), location_type(LOC), boolean_type(true), boolean_type(OUTPUT).

/*FLUENT STATUS*/
fun_fluent(status(P)) :- location_type(P).

causes_val(release(SRVC, ID, extinguishfire, _I, _E, [OUTPUT]), status(LOC), OUTPUT,
  assigned(SRVC, [workitem(extinguishfire, ID, [LOC], [ok])])) :-
  location_type(LOC), status_type(ok), status_type(OUTPUT).

causes_val(release(SRVC, ID, removedebris, _I, _E, [OUTPUT]), status(LOC), OUTPUT,
  assigned(SRVC, [workitem(removedebris, ID, [LOC], [ok])])) :-
  location_type(LOC), status_type(ok), status_type(OUTPUT).

causes_val(release(SRVC, ID, updatestatus, _I, _E, [OUTPUT]), status(LOC), OUTPUT,
  assigned(SRVC, [workitem(updatestatus, ID, [LOC], [ok])])) :-
  location_type(LOC), status_type(ok), status_type(OUTPUT).

fun_fluent(status_exp(P)) :- location_type(P).

causes_val(release(SRVC, ID, extinguishfire, _I, _E, [OUTPUT]), status_exp(LOC), ok,
  and(assigned(SRVC, [workitem(extinguishfire, ID, [LOC], [ok])]), neg(adapting))) :-
  location_type(LOC), status_type(ok), status_type(OUTPUT).

causes_val(release(SRVC, ID, removedebris, _I, _E, [OUTPUT]), status_exp(LOC), ok,
  and(assigned(SRVC, [workitem(removedebris, ID, [LOC], [ok])]), neg(adapting))) :-
  location_type(LOC), status_type(ok), status_type(OUTPUT).

causes_val(release(SRVC, ID, updatestatus, _I, _E, [OUTPUT]), status_exp(LOC), ok,
  and(assigned(SRVC, [workitem(updatestatus, ID, [LOC], [ok])]), neg(adapting))) :-
  location_type(LOC), status_type(ok), status_type(OUTPUT).

/*FLUENT BATTERYLEVEL*/
fun_fluent(batteryLevel(SRVC)) :- service(SRVC).

causes_val(release(SRVC, ID, move, _I, _E, [OUTPUT]), batteryLevel(SRVC), L,
  and(assigned(SRVC, [workitem(move, ID, [FROM, TO], [TO])]),
    L is batteryLevel(SRVC)-moveStep)) :-
  service(SRVC), location_type(OUTPUT), location_type(FROM), location_type(TO).

causes_val(release(SRVC, ID, removedebris, _I, _E, [OUTPUT]), batteryLevel(SRVC), L,
  and(assigned(SRVC, [workitem(removedebris, ID, [LOC], [ok])]),
    L is batteryLevel(SRVC)-debrisStep)) :-
  service(SRVC), location_type(LOC), status_type(ok), status_type(OUTPUT).

causes_val(release(SRVC, ID, chargebattery, _I, _E, []), batteryLevel(RB), L,
  and(assigned(SRVC, [workitem(chargebattery, ID, [RB], [])]),

```

```

                L is batteryLevel(RB)+batteryRecharging)) :-
    service(SRVC),service(RB).

/*FLUENT PHOTO TAKEN*/
fun_fluent(photoTaken(P)) :- location_type(P).

causes_val(release(SRVC,ID,takephoto,_I,_E,[OUTPUT]),photoTaken(LOC),OUTPUT,
    assigned(SRVC,[workitem(takephoto,ID,[LOC],[true])])) :-
    service(SRVC),id(ID),location_type(LOC),boolean_type(true),boolean_type(OUTPUT).

fun_fluent(photoTaken_exp(P)) :- location_type(P).

causes_val(release(SRVC,ID,takephoto,_I,_E,[OUTPUT]),photoTaken_exp(LOC),true,
    and(neg(adapting),assigned(SRVC,[workitem(takephoto,ID,[LOC],[true])])) :-
    service(SRVC),id(ID),location_type(LOC),boolean_type(true),boolean_type(OUTPUT).

/*FLUENT GENERAL BATTERY*/
fun_fluent(generalBattery).

causes_val(release(SRVC,ID,chargebattery,_I,_E,[ ]),generalBattery,L,
    and(assigned(SRVC,[workitem(chargebattery,ID,[RB],[ ])]),L is generalBattery-batteryRecharging)) :-
    service(SRVC),id(ID),service(RB).

/*FLUENT BATTERY RECHARGING*/
fun_fluent(batteryRecharging).

/*FLUENT MOVE STEP*/
fun_fluent(moveStep).

/*FLUENT DEBRIS STEP*/
fun_fluent(debrisStep).

/*ABBREVIATIONS*/

proc(isConnected(SRVC),
    and(provides(SRVC,movement),
    or(covered(at(SRVC)),
        some(rb,and(service(rb),and(provides(rb,battery),
            and(or(neigh(at(SRVC),atRobot(rb)),
                at(SRVC)=atRobot(rb)),isRobotConnected(rb))))))
    )).

proc(isRobotConnected(RBT),
    and(service(RBT),and(provides(RBT,battery),
        or(covered(atRobot(RBT)),
            or(some(loc,and(location_type(loc),and(covered(loc),neigh(atRobot(RBT),loc)))),
                or(some(rb,and(service(rb),
                    and(provides(rb,battery),and(neg(rb=RBT),and(neigh(atRobot(RBT),atRobot(rb))),
                or(covered(atRobot(rb)),
                    some(lc2,and(location_type(lc2),and(covered(lc2),neigh(atRobot(rb),lc2))))))))))
                    some(rb,and(service(rb),and(provides(rb,battery),and(neg(rb=RBT),
                        some(lc,and(location_type(lc),and(neigh(atRobot(RBT),lc),and(neigh(atRobot(rb),lc),
                or(covered(atRobot(rb)),
                    some(lc2,and(location_type(lc2),and(covered(lc2),neigh(atRobot(rb),lc2))))
                ))))))))
    )))))))

```

```

    )
  )
)
).

/* INITIAL STATE */

initially(reserved(SRVC, ID, TASK), false) :- task(TASK), service(SRVC), id(ID).
initially(assigned(SRVC, ID, TASK), false) :- task(TASK), service(SRVC), id(ID).
initially(finished, false).
initially(adapting, false).
initially(built_in_adaptation, false).

initially(free(act1), true).
initially(free(act2), true).
initially(free(act3), false).
initially(free(act4), false).
initially(free(rb1), true).
initially(free(rb2), true).

initially(batteryRecharging, 10).
initially(moveStep, 2).
initially(debrisStep, 3).
initially(generalBattery, 30).

initially(at(act1), loc00).
initially(at_exp(act1), loc00).
initially(at(act2), loc00).
initially(at_exp(act2), loc00).
initially(at(act3), loc00).
initially(at_exp(act3), loc00).
initially(at(act4), loc00).
initially(at_exp(act4), loc00).

initially(atRobot(rb1), loc00).
initially(atRobot(rb2), loc00).

initially(evacuated(loc00), false).
initially(evacuated_exp(loc00), false).
initially(evacuated(loc01), false).
initially(evacuated_exp(loc01), false).
initially(evacuated(loc02), false).
initially(evacuated_exp(loc02), false).
initially(evacuated(loc03), false).
initially(evacuated_exp(loc03), false).
initially(evacuated(loc10), false).
initially(evacuated_exp(loc10), false).
initially(evacuated(loc11), false).
initially(evacuated_exp(loc11), false).
initially(evacuated(loc13), false).
initially(evacuated_exp(loc13), false).
initially(evacuated(loc20), false).
initially(evacuated_exp(loc20), false).
initially(evacuated(loc23), false).
initially(evacuated_exp(loc23), false).
initially(evacuated(loc30), false).

```

```
initially(evacuated_exp(loc30),false).
initially(evacuated(loc31),false).
initially(evacuated_exp(loc31),false).
initially(evacuated(loc32),false).
initially(evacuated_exp(loc32),false).
initially(evacuated(loc33),false).
initially(evacuated_exp(loc33),false).
```

```
initially(photoTaken(loc00),false).
initially(photoTaken_exp(loc00),false).
initially(photoTaken(loc01),false).
initially(photoTaken_exp(loc01),false).
initially(photoTaken(loc02),false).
initially(photoTaken_exp(loc02),false).
initially(photoTaken(loc03),false).
initially(photoTaken_exp(loc03),false).
initially(photoTaken(loc10),false).
initially(photoTaken_exp(loc10),false).
initially(photoTaken(loc11),false).
initially(photoTaken_exp(loc11),false).
initially(photoTaken(loc13),false).
initially(photoTaken_exp(loc13),false).
initially(photoTaken(loc20),false).
initially(photoTaken_exp(loc20),false).
initially(photoTaken(loc23),false).
initially(photoTaken_exp(loc23),false).
initially(photoTaken(loc30),false).
initially(photoTaken_exp(loc30),false).
initially(photoTaken(loc31),false).
initially(photoTaken_exp(loc31),false).
initially(photoTaken(loc32),false).
initially(photoTaken_exp(loc32),false).
initially(photoTaken(loc33),false).
initially(photoTaken_exp(loc33),false).
```

```
initially(status(loc00),ok).
initially(status_exp(loc00),ok).
initially(status(loc01),ok).
initially(status_exp(loc01),ok).
initially(status(loc02),ok).
initially(status_exp(loc02),ok).
initially(status(loc03),ok).
initially(status_exp(loc03),ok).
initially(status(loc10),ok).
initially(status_exp(loc10),ok).
initially(status(loc11),ok).
initially(status_exp(loc11),ok).
initially(status(loc13),ok).
initially(status_exp(loc13),ok).
initially(status(loc20),ok).
initially(status_exp(loc20),ok).
initially(status(loc23),ok).
initially(status_exp(loc23),ok).
initially(status(loc30),ok).
initially(status_exp(loc30),ok).
initially(status(loc31),ok).
initially(status_exp(loc31),ok).
```

```

initially(status(loc32),ok).
initially(status_exp(loc32),ok).
initially(status(loc33),ok).
initially(status_exp(loc33),ok).

```

```

initially(batteryLevel(act1),0).
initially(batteryLevel(act2),0).
initially(batteryLevel(act3),0).
initially(batteryLevel(act4),0).
initially(batteryLevel(rb1),15).
initially(batteryLevel(rb2),15).

```

```

initially(covered(loc00),true).
initially(covered(loc10),true).
initially(covered(loc20),true).
initially(covered(loc11),true).
initially(covered(loc01),true).
initially(covered(loc02),true).

```

```

initially(neigh(loc00,loc10),true).
initially(neigh(loc00,loc11),true).
initially(neigh(loc00,loc01),true).
initially(neigh(loc11,loc10),true).
initially(neigh(loc11,loc01),true).
initially(neigh(loc11,loc00),true).
initially(neigh(loc11,loc20),true).
initially(neigh(loc11,loc02),true).
initially(neigh(loc10,loc20),true).
initially(neigh(loc10,loc00),true).
initially(neigh(loc10,loc11),true).
initially(neigh(loc10,loc01),true).
initially(neigh(loc01,loc02),true).
initially(neigh(loc01,loc11),true).
initially(neigh(loc01,loc10),true).
initially(neigh(loc01,loc00),true).
initially(neigh(loc02,loc03),true).
initially(neigh(loc02,loc13),true).
initially(neigh(loc02,loc01),true).
initially(neigh(loc02,loc11),true).
initially(neigh(loc03,loc02),true).
initially(neigh(loc03,loc13),true).
initially(neigh(loc13,loc03),true).
initially(neigh(loc13,loc23),true).
initially(neigh(loc13,loc02),true).
initially(neigh(loc23,loc13),true).
initially(neigh(loc23,loc33),true).
initially(neigh(loc23,loc32),true).
initially(neigh(loc33,loc23),true).
initially(neigh(loc33,loc32),true).
initially(neigh(loc32,loc33),true).
initially(neigh(loc32,loc23),true).
initially(neigh(loc32,loc31),true).
initially(neigh(loc31,loc32),true).
initially(neigh(loc31,loc20),true).
initially(neigh(loc31,loc30),true).
initially(neigh(loc30,loc31),true).
initially(neigh(loc30,loc20),true).

```

```

initially(neigh(loc20,loc30),true).
initially(neigh(loc20,loc31),true).
initially(neigh(loc20,loc10),true).
initially(neigh(loc20,loc11),true).

/* EXOGENOUS EVENTS USED FOR SWITCHING A TASK FROM A STATE TO ANOTHER */
prim_action(A) :- exog_action(A).
poss(A,true) :- exog_action(A).

exog_action(readyToStart(SRVC,ID,TASK)) :- task(TASK), service(SRVC), id(ID).
exog_action(finishedTask(SRVC,ID,TASK,_V)) :- task(TASK), service(SRVC), id(ID).

/* INTERNAL EXOGENOUS ACTION FOR LOADING THE RECOVERY PROCEDURE */
exog_action(planReady(_PLAN)).

/* DOMAIN-DEPENDENT EXOGENOUS EVENTS */

exog_action(photoLost(LOC)) :- location_type(LOC).
causes_val(photoLost(LOC),photoTaken(LOC),false,true).

exog_action(fireRisk(LOC)) :- location_type(LOC).
causes_val(fireRisk(LOC),status(LOC),fire,true).

exog_action(rockSlide(LOC)) :- location_type(LOC).
causes_val(rockSlide(LOC),status(LOC),debris,true).

/* INDIGOLOG INTERNAL ACTIONS */

prim_action(initPMS).
poss(initPMS, true).

prim_action(endPMS).
poss(endPMS, true).

prim_action(finish).
poss(finish,true).
rel_fluent(finished).
causes_val(finish,finished,true,true).

rel_fluent(realityChanged).
initially(realityChanged,false).
causes_val(release(_,_),realityChanged,true,neg(adapting)).

prim_action(resetReality).
poss(resetReality,true).
causes_val(resetReality,realityChanged,false,true).

prim_action(adaptFinish).
poss(adaptFinish,true).

prim_action(adaptStart).
poss(adaptStart,true).

prim_action(adaptFound).
poss(adaptFound,true).

prim_action(adaptNotFound).

```

```

poss(adaptNotFound,true).

/*COMMANDS FOR INVOKING THE PLANNER*/
prim_action(invokePlanner).
poss(invokePlanner,true).

rel_fluent(adaptationPlanReady).
initially(adaptationPlanReady,true).
causes_val(invokePlanner,adaptationPlanReady,false,true).
causes_val(planReady(_P),adaptationPlanReady,true,true).

fun_fluent(recoveryPlan).
initially(recoveryPlan,[]).
causes_val(planReady(P),recoveryPlan,P,true).

prim_action(stop).
poss(stop,false).

rel_fluent(adapting).
causes_val(adaptStart,adapting,true,true).
causes_val(adaptFound,adapting,false,true).

/* DEFINITION OF THE RELEVANT FLUENT */
proc(relevant,or(some(srv, and(service(srv), and(provides(srv,movement),
    or(neg(at(srv)=at_exp(srv)),neg(isConnected(srv)))))),
    some(poi, and(location_type(poi), or(neg(evacuated_exp(poi)=evacuated(poi)),
    or(neg(photoTaken_exp(poi)=photoTaken(poi)),neg(status_exp(poi)=status(poi)))))))).

/* THE MONITOR PROCEDURE */
proc(monitor,[?(writeln('Monitor')),ndet([?(neg(relevant)),?(writeln('NonRelevant'))],
    [?(relevant),?(writeln('Relevant'))],
    [adaptStart,createPlanningProblem,?(report_message(user, 'About to adapt...')),
        if(built_in_adaptation,pconc([adaptingProgram,adaptFinish],
            while(adapting,[?(writeln('waiting')),wait])),
    [invokePlanner,manageRecoveryProcess(recoveryPlan),adaptFound,adaptFinish]
    ])],
    resetReality]).

/* THE BUILT-IN ADAPTATION PROCEDURE */
proc(adaptingProgram,searchn([?(true),adapt,[adaptFound,
    ?(report_message(user, 'Adaptation program found!'))]],
    [assumptions([[assign(N,D,T,I,E)],readyToStart(N,D,T)],
        [start(N,D,T),finishedTask(N,D,T,E)]])).

proc(adapt,plans(0,10)).

proc(plans(M,N),[if(M=N,[adaptNotFound],[?(M<(N+1))],
    ndet([actionSequence(M),?(neg(relevant))],
        [?(SUCCM is M+1),plans(SUCCM,N)]))]).

proc(actionSequence(N),ndet([?(N=0)],[?(N>0),
    pi([n,t,i,e],[?(and(service(n),and(free(n),capable(n,[workitem(t,id_adapt,i,e)])))]),
    assign(n,id_adapt,t,i,e),
    start(n,id_adapt,t),
    ackCompl(n,id_adapt,t),
    release(n,id_adapt,t,i,e),
    ?(PRECN is N-1), actionSequence(PRECN)]))

```



```

proc(createPlanningProblem,[writeFile('Planners/LPG-TD/problem.pddl')]).

proc(writeFile(File),[?(open(File, write, Stream)),
?(writeln(Stream,('define (problem EM1) (:domain Derailment')))),
?(writeln(Stream,('(:objects')))),
?(writeln(Stream,('act1 - service'))),
?(writeln(Stream,('act2 - service'))),
?(writeln(Stream,('act3 - service'))),
?(writeln(Stream,('act4 - service'))),
?(writeln(Stream,('rb1 - service'))),
?(writeln(Stream,('rb2 - service'))),
?(writeln(Stream,('ok - status_type'))),
?(writeln(Stream,('fire - status_type'))),
?(writeln(Stream,('debris - status_type'))),
?(writeln(Stream,('movement - capability'))),
?(writeln(Stream,('hatchet - capability'))),
?(writeln(Stream,('camera - capability'))),
?(writeln(Stream,('gprs - capability'))),
?(writeln(Stream,('extinguisher - capability'))),
?(writeln(Stream,('battery - capability'))),
?(writeln(Stream,('digger - capability'))),
?(writeln(Stream,('powerpack - capability'))),
?(writeln(Stream,('loc00 - location_type'))),
?(writeln(Stream,('loc01 - location_type'))),
?(writeln(Stream,('loc02 - location_type'))),
?(writeln(Stream,('loc03 - location_type'))),
?(writeln(Stream,('loc10 - location_type'))),
?(writeln(Stream,('loc11 - location_type'))),
?(writeln(Stream,('loc13 - location_type'))),
?(writeln(Stream,('loc20 - location_type'))),
?(writeln(Stream,('loc23 - location_type'))),
?(writeln(Stream,('loc30 - location_type'))),
?(writeln(Stream,('loc31 - location_type'))),
?(writeln(Stream,('loc32 - location_type'))),
?(writeln(Stream,('loc33 - location_type'))),
?(writeln(Stream,(''))),
?(writeln(Stream,('(:init'))),
if(free(act1)=true,?(writeln(Stream,('free act1'))),[]),
if(free(act2)=true,?(writeln(Stream,('free act2'))),[]),
if(free(act3)=true,?(writeln(Stream,('free act3'))),[]),
if(free(act4)=true,?(writeln(Stream,('free act4'))),[]),
if(free(rb1)=true,?(writeln(Stream,('free rb1'))),[]),
if(free(rb2)=true,?(writeln(Stream,('free rb2'))),[]),
?(writeln(Stream,('provides act1 movement'))),
?(writeln(Stream,('provides act1 gprs'))),
?(writeln(Stream,('provides act1 camera'))),
?(writeln(Stream,('provides act1 extinguisher'))),
?(writeln(Stream,('provides act2 movement'))),
?(writeln(Stream,('provides act2 gprs'))),
?(writeln(Stream,('provides act2 camera'))),
?(writeln(Stream,('provides act2 hatchet'))),
?(writeln(Stream,('provides act3 movement'))),
?(writeln(Stream,('provides act3 gprs'))),
?(writeln(Stream,('provides act3 camera'))),
?(writeln(Stream,('provides act3 hatchet'))),
?(writeln(Stream,('provides act4 movement'))),
?(writeln(Stream,('provides act4 powerpack'))),

```

```
?(writeln(Stream,(' provides act4 gprs'))),
?(writeln(Stream,(' provides rb1 battery'))),
?(writeln(Stream,(' provides rb1 digger'))),
?(writeln(Stream,(' provides rb2 battery'))),
?(writeln(Stream,(' provides rb2 digger'))),
?(writeln(Stream,(' covered loc00'))),
?(writeln(Stream,(' covered loc10'))),
?(writeln(Stream,(' covered loc20'))),
?(writeln(Stream,(' covered loc01'))),
?(writeln(Stream,(' covered loc11'))),
?(writeln(Stream,(' covered loc02'))),
?(writeln(Stream,(' neigh loc00 loc10'))),
?(writeln(Stream,(' neigh loc00 loc11'))),
?(writeln(Stream,(' neigh loc00 loc01'))),
?(writeln(Stream,(' neigh loc11 loc10'))),
?(writeln(Stream,(' neigh loc11 loc01'))),
?(writeln(Stream,(' neigh loc11 loc00'))),
?(writeln(Stream,(' neigh loc11 loc20'))),
?(writeln(Stream,(' neigh loc11 loc02'))),
?(writeln(Stream,(' neigh loc10 loc20'))),
?(writeln(Stream,(' neigh loc10 loc00'))),
?(writeln(Stream,(' neigh loc10 loc11'))),
?(writeln(Stream,(' neigh loc10 loc01'))),
?(writeln(Stream,(' neigh loc01 loc02'))),
?(writeln(Stream,(' neigh loc01 loc11'))),
?(writeln(Stream,(' neigh loc01 loc10'))),
?(writeln(Stream,(' neigh loc01 loc00'))),
?(writeln(Stream,(' neigh loc02 loc03'))),
?(writeln(Stream,(' neigh loc02 loc13'))),
?(writeln(Stream,(' neigh loc02 loc01'))),
?(writeln(Stream,(' neigh loc02 loc11'))),
?(writeln(Stream,(' neigh loc03 loc02'))),
?(writeln(Stream,(' neigh loc03 loc13'))),
?(writeln(Stream,(' neigh loc13 loc03'))),
?(writeln(Stream,(' neigh loc13 loc23'))),
?(writeln(Stream,(' neigh loc13 loc02'))),
?(writeln(Stream,(' neigh loc23 loc13'))),
?(writeln(Stream,(' neigh loc23 loc33'))),
?(writeln(Stream,(' neigh loc23 loc32'))),
?(writeln(Stream,(' neigh loc33 loc23'))),
?(writeln(Stream,(' neigh loc33 loc32'))),
?(writeln(Stream,(' neigh loc32 loc33'))),
?(writeln(Stream,(' neigh loc32 loc23'))),
?(writeln(Stream,(' neigh loc32 loc31'))),
?(writeln(Stream,(' neigh loc31 loc32'))),
?(writeln(Stream,(' neigh loc31 loc20'))),
?(writeln(Stream,(' neigh loc31 loc30'))),
?(writeln(Stream,(' neigh loc30 loc31'))),
?(writeln(Stream,(' neigh loc30 loc20'))),
?(writeln(Stream,(' neigh loc20 loc30'))),
?(writeln(Stream,(' neigh loc20 loc31'))),
?(writeln(Stream,(' neigh loc20 loc10'))),
?(writeln(Stream,(' neigh loc20 loc11'))),
?(write(Stream,(' ')),?(write(Stream,('at act1 '))),
?(write(Stream,(at(act1))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('at act2 '))),
?(write(Stream,(at(act2))),?(writeln(Stream,(' '))),
```

```

?(write(Stream,(' ')),?(write(Stream,('at act3 '))),
?(write(Stream,(at(act3))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('at act4 '))),
?(write(Stream,(at(act4))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('atRobot rb1 '))),
?(write(Stream,(atRobot(rb1))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('atRobot rb2 '))),
?(write(Stream,(atRobot(rb2))),?(writeln(Stream,(' '))),
?(write(Stream,('= ')),?(write(Stream,('batteryLevel rb1 '))),
?(write(Stream,(batteryLevel(rb1))),?(writeln(Stream,(' '))),
?(write(Stream,('= ')),?(write(Stream,('batteryLevel rb2 '))),
?(write(Stream,(batteryLevel(rb2))),?(writeln(Stream,(' '))),
?(write(Stream,('= ')),?(write(Stream,('batteryRecharging '))),
?(write(Stream,(batteryRecharging)),?(writeln(Stream,(' '))),
?(write(Stream,('= ')),?(write(Stream,('generalBattery '))),
?(write(Stream,(generalBattery)),?(writeln(Stream,(' '))),
?(write(Stream,('= ')),?(write(Stream,('debrisStep '))),
?(write(Stream,(debrisStep)),?(writeln(Stream,(' '))),
?(write(Stream,('= ')),?(write(Stream,('moveStep '))),
?(write(Stream,(moveStep)),?(writeln(Stream,(' '))),
if(photoTaken(loc00)=true,?(writeln(Stream,('photoTaken loc00'))),[]),
if(photoTaken(loc01)=true,?(writeln(Stream,('photoTaken loc01'))),[]),
if(photoTaken(loc02)=true,?(writeln(Stream,('photoTaken loc02'))),[]),
if(photoTaken(loc03)=true,?(writeln(Stream,('photoTaken loc03'))),[]),
if(photoTaken(loc10)=true,?(writeln(Stream,('photoTaken loc10'))),[]),
if(photoTaken(loc11)=true,?(writeln(Stream,('photoTaken loc11'))),[]),
if(photoTaken(loc13)=true,?(writeln(Stream,('photoTaken loc13'))),[]),
if(photoTaken(loc20)=true,?(writeln(Stream,('photoTaken loc20'))),[]),
if(photoTaken(loc23)=true,?(writeln(Stream,('photoTaken loc23'))),[]),
if(photoTaken(loc30)=true,?(writeln(Stream,('photoTaken loc30'))),[]),
if(photoTaken(loc31)=true,?(writeln(Stream,('photoTaken loc31'))),[]),
if(photoTaken(loc32)=true,?(writeln(Stream,('photoTaken loc32'))),[]),
if(photoTaken(loc33)=true,?(writeln(Stream,('photoTaken loc33'))),[]),
?(write(Stream,(' ')),?(write(Stream,('status loc00 '))),
?(write(Stream,(status(loc00))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('status loc01 '))),
?(write(Stream,(status(loc01))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('status loc02 '))),
?(write(Stream,(status(loc02))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('status loc03 '))),
?(write(Stream,(status(loc03))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('status loc10 '))),
?(write(Stream,(status(loc10))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('status loc11 '))),
?(write(Stream,(status(loc11))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('status loc13 '))),
?(write(Stream,(status(loc13))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('status loc20 '))),
?(write(Stream,(status(loc20))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('status loc23 '))),
?(write(Stream,(status(loc23))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('status loc30 '))),
?(write(Stream,(status(loc30))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('status loc31 '))),
?(write(Stream,(status(loc31))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('status loc32 '))),
?(write(Stream,(status(loc32))),?(writeln(Stream,(' '))),

```



```

?(write(Stream,(' ')),?(write(Stream,('status loc33 '))),
?(write(Stream,(status(loc33))),?(writeln(Stream,(' '))),
if(evacuated(loc00)=true,?(writeln(Stream,('evacuated loc00'))),[]),
if(evacuated(loc01)=true,?(writeln(Stream,('evacuated loc01'))),[]),
if(evacuated(loc02)=true,?(writeln(Stream,('evacuated loc02'))),[]),
if(evacuated(loc03)=true,?(writeln(Stream,('evacuated loc03'))),[]),
if(evacuated(loc10)=true,?(writeln(Stream,('evacuated loc10'))),[]),
if(evacuated(loc11)=true,?(writeln(Stream,('evacuated loc11'))),[]),
if(evacuated(loc13)=true,?(writeln(Stream,('evacuated loc13'))),[]),
if(evacuated(loc20)=true,?(writeln(Stream,('evacuated loc20'))),[]),
if(evacuated(loc23)=true,?(writeln(Stream,('evacuated loc23'))),[]),
if(evacuated(loc30)=true,?(writeln(Stream,('evacuated loc30'))),[]),
if(evacuated(loc31)=true,?(writeln(Stream,('evacuated loc31'))),[]),
if(evacuated(loc32)=true,?(writeln(Stream,('evacuated loc32'))),[]),
if(evacuated(loc33)=true,?(writeln(Stream,('evacuated loc33'))),[]),
?(writeln(Stream,(' '))),
?(writeln(Stream,(':goal'))),
?(writeln(Stream,('and'))),
?(write(Stream,(' ')),?(write(Stream,('at act1 '))),
?(write(Stream,(at_exp(act1))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('at act2 '))),
?(write(Stream,(at_exp(act2))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('at act3 '))),
?(write(Stream,(at_exp(act3))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('at act4 '))),
?(write(Stream,(at_exp(act4))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('status loc00 '))),
?(write(Stream,(status_exp(loc00))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('status loc01 '))),
?(write(Stream,(status_exp(loc01))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('status loc02 '))),
?(write(Stream,(status_exp(loc02))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('status loc03 '))),
?(write(Stream,(status_exp(loc03))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('status loc10 '))),
?(write(Stream,(status_exp(loc10))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('status loc11 '))),
?(write(Stream,(status_exp(loc11))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('status loc13 '))),
?(write(Stream,(status_exp(loc13))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('status loc20 '))),
?(write(Stream,(status_exp(loc20))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('status loc23 '))),
?(write(Stream,(status_exp(loc23))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('status loc30 '))),
?(write(Stream,(status_exp(loc30))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('status loc31 '))),
?(write(Stream,(status_exp(loc31))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('status loc32 '))),
?(write(Stream,(status_exp(loc32))),?(writeln(Stream,(' '))),
?(write(Stream,(' ')),?(write(Stream,('status loc33 '))),
?(write(Stream,(status_exp(loc33))),?(writeln(Stream,(' '))),
if(evacuated_exp(loc00)=true,?(writeln(Stream,('evacuated loc00'))),[]),
if(evacuated_exp(loc01)=true,?(writeln(Stream,('evacuated loc01'))),[]),
if(evacuated_exp(loc02)=true,?(writeln(Stream,('evacuated loc02'))),[]),
if(evacuated_exp(loc03)=true,?(writeln(Stream,('evacuated loc03'))),[]),
if(evacuated_exp(loc10)=true,?(writeln(Stream,('evacuated loc10'))),[]),

```

```

if(evacuated_exp(loc11)=true,?(writeln(Stream,('evacuated loc11'))),[]),
if(evacuated_exp(loc13)=true,?(writeln(Stream,('evacuated loc13'))),[]),
if(evacuated_exp(loc20)=true,?(writeln(Stream,('evacuated loc20'))),[]),
if(evacuated_exp(loc23)=true,?(writeln(Stream,('evacuated loc23'))),[]),
if(evacuated_exp(loc30)=true,?(writeln(Stream,('evacuated loc30'))),[]),
if(evacuated_exp(loc31)=true,?(writeln(Stream,('evacuated loc31'))),[]),
if(evacuated_exp(loc32)=true,?(writeln(Stream,('evacuated loc32'))),[]),
if(evacuated_exp(loc33)=true,?(writeln(Stream,('evacuated loc33'))),[]),
if(photoTaken_exp(loc00)=true,?(writeln(Stream,('photoTaken loc00'))),[]),
if(photoTaken_exp(loc01)=true,?(writeln(Stream,('photoTaken loc01'))),[]),
if(photoTaken_exp(loc02)=true,?(writeln(Stream,('photoTaken loc02'))),[]),
if(photoTaken_exp(loc03)=true,?(writeln(Stream,('photoTaken loc03'))),[]),
if(photoTaken_exp(loc10)=true,?(writeln(Stream,('photoTaken loc10'))),[]),
if(photoTaken_exp(loc11)=true,?(writeln(Stream,('photoTaken loc11'))),[]),
if(photoTaken_exp(loc13)=true,?(writeln(Stream,('photoTaken loc13'))),[]),
if(photoTaken_exp(loc20)=true,?(writeln(Stream,('photoTaken loc20'))),[]),
if(photoTaken_exp(loc23)=true,?(writeln(Stream,('photoTaken loc23'))),[]),
if(photoTaken_exp(loc30)=true,?(writeln(Stream,('photoTaken loc30'))),[]),
if(photoTaken_exp(loc31)=true,?(writeln(Stream,('photoTaken loc31'))),[]),
if(photoTaken_exp(loc32)=true,?(writeln(Stream,('photoTaken loc32'))),[]),
if(photoTaken_exp(loc33)=true,?(writeln(Stream,('photoTaken loc33'))),[]),
?(writeln(Stream,('isConnected act1'))),
?(writeln(Stream,('isConnected act2'))),
?(writeln(Stream,('isConnected act3'))),
?(writeln(Stream,('isConnected act4'))),
?(writeln(Stream,(''))),
?(writeln(Stream,('(:metric minimize (total-time))')),
?(writeln(Stream,(''))),
?(nl(Stream),?(close(Stream))]).

```

PDDL Planning Domain

```

(define (domain Derailment)
  (:requirements :derived-predicates :typing :fluents :equality)
  (:types service capability location_type status_type)

  (:predicates
    (free ?x - service)
    (provides ?x - service ?c - capability)
    (neigh ?y1 - location_type ?y2 - location_type)
    (covered ?y - location_type)
    (at ?x - service ?y - location_type)
    (atRobot ?x - service ?y - location_type)
    (status ?y - location_type ?s - status_type)
    (evacuated ?y - location_type)
    (photoTaken ?l - location_type)
    (isRobotConnected ?x - service)
    (isConnected ?x - service)
  )

  (:functions
    (batteryLevel ?x - service)
    (moveStep)
    (debrisStep)
    (generalBattery)
  )

```

```

(batteryRecharging)
)

(:action go
:parameters (?x - service ?from - location_type ?to - location_type)
:precondition (and (provides ?x movement) (free ?x) (at ?x ?from) (isConnected ?x))
:effect (and (not (at ?x ?from)) (at ?x ?to))
)

(:action move
:parameters (?x - service ?from - location_type ?to - location_type)
:precondition (and (provides ?x battery) (free ?x) (atRobot ?x ?from) (>= (batteryLevel ?x)
(moveStep)) (isRobotConnected ?x))
:effect (and (not (atRobot ?x ?from)) (atRobot ?x ?to) (decrease (batteryLevel ?x) (moveStep)))
)

(:action extinguishfire
:parameters (?x - service ?y - location_type)
:precondition (and (provides ?x extinguisher) (free ?x) (at ?x ?y)
(status ?y fire) (isConnected ?x))
:effect (status ?y ok)
)

(:action evacuate
:parameters (?x - service ?y - location_type)
:precondition (and (provides ?x hatchet) (free ?x) (at ?x ?y) (status ?y ok)
(not (evacuated ?y)) (isConnected ?x))
:effect (evacuated ?y)
)

(:action removedebris
:parameters (?x - service ?y - location_type)
:precondition (and (provides ?x digger) (free ?x) (atRobot ?x ?y) (status ?y debris)
(>= (batteryLevel ?x) (debrisStep)) (isRobotConnected ?x))
:effect (and (status ?y ok) (decrease (batteryLevel ?x) (debrisStep)))
)

(:action chargebattery
:parameters (?x1 - service ?x2 - service ?y - location_type)
:precondition (and (provides ?x1 powerpack) (free ?x1) (at ?x1 ?y) (provides ?x2 battery)
(atRobot ?x2 ?y) (>= (generalBattery) (batteryRecharging)) (isConnected ?x1))
:effect (and (decrease (generalBattery) (batteryRecharging))
(increase (batteryLevel ?x2) (batteryRecharging)))
)

(:action updatestatus
:parameters (?x - service ?y - location_type)
:precondition (and (provides ?x gprs) (free ?x) (at ?x ?y) (status ?y ok) (isConnected ?x))
:effect (status ?y ok)
)

(:action takephoto
:parameters (?x - service ?y - location_type)
:precondition (and (provides ?x camera) (free ?x) (at ?x ?y) (isConnected ?x))
:effect (photoTaken ?y)
)

```

```

(:derived (isConnected ?x - service)
  (and (provides ?x movement)
    (or (exists (?y - location_type) (and (at ?x ?y) (covered ?y)))
      (exists (?r - service ?y - location_type ?z - location_type)
        (and (provides ?r battery) (at ?x ?y) (atRobot ?r ?z)
          (or (neigh ?y ?z) (= ?y ?z)) (isRobotConnected ?r)))
      )
    )
  )
)

(:derived (isRobotConnected ?x - service) (and (provides ?x battery)
  (or (exists (?y - location_type) (and (atRobot ?x ?y) (covered ?y)))
    (exists (?y - location_type ?z - location_type)
      (and (atRobot ?x ?y) (covered ?z) (neigh ?y ?z)))
    (exists (?r - service ?y - location_type ?z - location_type ?k - location_type)
      (and (provides ?r battery) (not (= ?x ?r))
        (atRobot ?x ?y) (atRobot ?r ?z) (neigh ?y ?z) (or (covered ?z)
          (and (covered ?k) (neigh ?k ?z)))))
      (exists (?r - service ?y - location_type ?z - location_type
        ?k - location_type ?c - location_type)
        (and (provides ?r battery)
          (not (= ?x ?r)) (atRobot ?x ?y) (atRobot ?r ?z) (neigh ?y ?k) (neigh ?z ?k)
            (or (covered ?z) (and (covered ?k) (neigh ?k ?z))) )
        )
      )
    )
  )
)
)
)
)

```

PDDL Planning Problem

```

(define (problem EM1) (:domain Derailment)
  (:objects
    act1 - service
    act2 - service
    act3 - service
    act4 - service
    rb1 - service
    rb2 - service
    ok - status_type
    fire - status_type
    debris - status_type
    movement - capability
    hatchet - capability
    camera - capability
    gprs - capability
    extinguisher - capability
    battery - capability
    digger - capability
    powerpack - capability
    loc00 - location_type
    loc01 - location_type
    loc02 - location_type
    loc03 - location_type
    loc10 - location_type
    loc11 - location_type
    loc13 - location_type
  )
)

```

```
loc20 - location_type
loc23 - location_type
loc30 - location_type
loc31 - location_type
loc32 - location_type
loc33 - location_type
)
(:init
  (free act1)
  (free act2)
  (free act3)
  (free act4)
  (free rb1)
  (free rb2)
  (provides act1 movement)
  (provides act1 gprs)
  (provides act1 camera)
  (provides act1 extinguisher)
  (provides act2 movement)
  (provides act2 gprs)
  (provides act2 camera)
  (provides act2 hatchet)
  (provides act3 movement)
  (provides act3 gprs)
  (provides act3 camera)
  (provides act3 hatchet)
  (provides act4 movement)
  (provides act4 powerpack)
  (provides act4 gprs)
  (provides rb1 battery)
  (provides rb1 digger)
  (provides rb2 battery)
  (provides rb2 digger)
  (covered loc00)
  (covered loc10)
  (covered loc20)
  (covered loc01)
  (covered loc11)
  (covered loc02)
  (neigh loc00 loc10)
  (neigh loc00 loc11)
  (neigh loc00 loc01)
  (neigh loc11 loc10)
  (neigh loc11 loc01)
  (neigh loc11 loc00)
  (neigh loc11 loc20)
  (neigh loc11 loc02)
  (neigh loc10 loc20)
  (neigh loc10 loc00)
  (neigh loc10 loc11)
  (neigh loc10 loc01)
  (neigh loc01 loc02)
  (neigh loc01 loc11)
  (neigh loc01 loc10)
  (neigh loc01 loc00)
  (neigh loc02 loc03)
  (neigh loc02 loc13)
```

```

(neigh loc02 loc01)
(neigh loc02 loc11)
(neigh loc03 loc02)
(neigh loc03 loc13)
(neigh loc13 loc03)
(neigh loc13 loc23)
(neigh loc13 loc02)
(neigh loc23 loc13)
(neigh loc23 loc33)
(neigh loc23 loc32)
(neigh loc33 loc23)
(neigh loc33 loc32)
(neigh loc32 loc33)
(neigh loc32 loc23)
(neigh loc32 loc31)
(neigh loc31 loc32)
(neigh loc31 loc20)
(neigh loc31 loc30)
(neigh loc30 loc31)
(neigh loc30 loc20)
(neigh loc20 loc30)
(neigh loc20 loc31)
(neigh loc20 loc10)
(neigh loc20 loc11)
(at act1 loc03)
(at act2 loc00)
(at act3 loc00)
(at act4 loc00)
(atRobot rb1 loc00)
(atRobot rb2 loc00)
(= (batteryLevel rb1) 15)
(= (batteryLevel rb2) 15)
(= (batteryRecharging) 10)
(= (generalBattery) 30)
(= (debrisStep) 3)
(= (moveStep) 2)
(status loc00 ok)
(status loc01 ok)
(status loc02 ok)
(status loc03 ok)
(status loc10 ok)
(status loc11 ok)
(status loc13 ok)
(status loc20 ok)
(status loc23 ok)
(status loc30 ok)
(status loc31 ok)
(status loc32 ok)
(status loc33 ok)
)
(:goal
(and
(at act1 loc33)
(at act2 loc00)
(at act3 loc00)
(at act4 loc00)
(status loc00 ok)

```

```
(status loc01 ok)
(status loc02 ok)
(status loc03 ok)
(status loc10 ok)
(status loc11 ok)
(status loc13 ok)
(status loc20 ok)
(status loc23 ok)
(status loc30 ok)
(status loc31 ok)
(status loc32 ok)
(status loc33 ok)
(isConnected act1)
(isConnected act2)
(isConnected act3)
(isConnected act4)
))
(:metric minimize (total-time))
)
```


Bibliography

- [1] ADAMS, M., TER HOFSTEDÉ, A. H. M., VAN DER AALST, W. M. P., AND EDMOND, D. Dynamic, extensible and context-aware exception handling for workflows. In *Proceedings of the 2007 OTM Confederated international conferences : CoopIS, DOA, ODBASE, GADA, and IS - Volume Part I*, OTM'07, pp. 95–112. Springer-Verlag (2007).
- [2] ADAMS, M. J. *Facilitating dynamic flexibility and exception handling for workflows*. Ph.D. thesis, Queensland University of Technology Brisbane, Australia (2007).
- [3] ALER, R., BORRAJO, D., AND CAMACHO, D. A Knowledge-based Approach for Business Process Reengineering, SHAMASH. *Know.-Based Syst.*, **15** (2002).
- [4] BANDINELLI, S. C., FUGGETTA, A., AND GHEZZI, C. Software process model evolution in the spade environment. *Software Engineering, IEEE Transactions on*, **19** (1993).
- [5] BECKSTEIN, C. AND KLAUSNER, J. A Meta Level Architecture for Workflow Management. *Journal of Integrated Design and Process Science*, **3** (1999).
- [6] BERENS, P. The flower case-handling approach: beyond workflow management. *Process-Aware Information Systems*, (2005).
- [7] BPMI.ORG AND OMG. Business Process Modeling Notation - Final Specification Ver.2.0. <http://www.omg.org/spec/BPMN/2.0/PDF/> (2011).
- [8] BRACHMAN, R. AND LEVESQUE, H. *Knowledge Representation and Reasoning*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (2004). ISBN 1558609326.
- [9] BUCCHIARONE, A., PISTORE, M., RAIK, H., AND KAZHAMIKIN, R. Adaptation of service-based business processes by context-aware replanning. In *SOCA*, pp. 1–8 (2011).
- [10] CAPATA, A., MARRELLA, A., RUSSO, R., BORTENSCHLAGER, M., AND RIESER, H. A Geo-based Application for the Management of Mobile Actors during Crisis Situations. In *5th International Conference on Information Systems for Crisis Response and Management (ISCRAM 2008)* (2008).
- [11] CASATI, F. A discussion on approaches to handling exceptions in workflows. *ACM SIGGROUP Bulletin*, **20** (1999).

- [12] CASATI, F., CERI, S., PERNICI, B., POZZI, G., AND POZZI, G. Workflow evolution. pp. 211–238 (1998).
- [13] CASATI, F. AND CUGOLA, G. Error handling in process support systems. *Advances in Exception Handling Techniques*, (2001).
- [14] CATARCI, T., DE LEONI, M., MARRELLA, A., MECELLA, M., BORTENSCHLAGER, M., AND STEINMANN, R. The WORKPAD Project Experience: Improving the Disaster Response through Process Management and Geo Collaboration. In *7th International Conference on Information Systems for Crisis Response and Management (ISCRAM 2010)* (2010).
- [15] CATARCI, T., DE LEONI, M., MARRELLA, A., MECELLA, M., RUSSO, A., STEINMANN, R., AND BORTENSCHLAGER, M. WORKPAD: Process Management and Geo-Collaboration Help Disaster Response. *IJISCRAM*, **3** (2011), 32. doi:10.4018/jiscrm.2011010103.
- [16] CATARCI, T., ET AL. Pervasive Software Environments for Supporting Disaster Responses. *IEEE Internet Computing*, **12** (2008), 26. doi:10.1109/MIC.2008.18.
- [17] C.HAGEN AND G.ALONSO. Exception Handling in Workflow Management Systems. *IEEE Trans. Software Engineering*, **26** (2000).
- [18] CHIU, D. K. W., LI, Q., AND KARLAPALEM, K. A Logical Framework for Exception Handling in ADOBE Workflow Management System. In *12th International Conference on Advanced Information Systems Engineering, CAISE '00*, pp. 110–125. Springer-Verlag (2000).
- [19] COLES, A. J., COLES, A., FOX, M., AND LONG, D. Forward-Chaining Partial-Order Planning. In *ICAPS* (2010).
- [20] COSA GMBH. COSA BPM product description. Retrieved May 30, 2013, from: http://www.cosa.nl/docs/EN/COSA%20BPM%205.7%20Productdescription_en_new_K.pdf (2013).
- [21] COSSU, F., MARRELLA, A., MECELLA, M., RUSSO, A., BERTAZZONI, G., SUPPA, M., AND GRASSO, F. Improving Operational Support in Hospital Wards through Vocal Interfaces and Process-Awareness. In *25th IEEE International Symposium on Computer-Based Medical Systems (CBMS 2012)* (2012). doi:10.1109/CBMS.2012.6266329.
- [22] CUMBERLIDGE, M. *Business Process Management with JBoss jBPM: A Practical Guide for Business Analysts*. Packt Publishing (2007).
- [23] DADAM, P. AND RINDERLE, S. Workflow evolution. In *Encyclopedia of Database Systems*, pp. 3540–3544 (2009).
- [24] DE GIACOMO, G., LESPÉRANCE, Y., AND LEVESQUE, H. J. Congolog, a concurrent programming language based on the situation calculus. *Artificial Intelligence*, **121** (2000).

- [25] DE GIACOMO, G., LESPÉRANCE, Y., LEVESQUE, H., AND SARDINA, S. Indigolog: A high-level programming language for embedded reasoning agents. In *Multi-Agent Programming* (edited by A. El Fallah Seghrouchni, J. Dix, M. Dastani, and R. H. Bordini), pp. 31–72. Springer US (2009).
- [26] DE GIACOMO, G. AND LEVESQUE, H. J. An incremental interpreter for high-level programs with sensing. *Logical Foundations for Cognitive Agents*, (1999).
- [27] DE GIACOMO, G., REITER, R., AND SOUTCHANSKI, M. Execution Monitoring of High-Level Robot Programs. In *KR'98: Proceedings of the Sixth International Conference on Principles of Knowledge Representation and Reasoning*, pp. 453–465 (1998).
- [28] DE LEONI, M., MARRELLA, A., MECELLA, M., DE ROSA, F., POGGI, A., KREK, A., AND MANTI, F. Emergency Management: from User Requirements to a Flexible P2P Architecture. In *4th International Conference on Information Systems for Crisis Response and Management (ISCRAM 2007)* (2007).
- [29] DE LEONI, M., MARRELLA, A., MECELLA, M., AND SARDINA, S. SmartPM - Featuring Automatic Adaptation to Unplanned Exceptions. Tech. rep., Dipartimento di Informatica e Sistemistica ANTONIO RUBERTI, SAPIENZA - Università di Roma (2011).
- [30] DE LEONI, M., MARRELLA, A., MECELLA, M., VALENTINI, S., AND SARDIÑA, S. Coordinating Mobile Actors in Pervasive and Mobile Scenarios: An AI-Based Approach. In *17th Workshop on Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE 2008)*, pp. 82–87. IEEE (2008). doi: [10.1109/WETICE.2008.30](https://doi.org/10.1109/WETICE.2008.30).
- [31] DE LEONI, M., MARRELLA, A., AND RUSSO, A. Process-Aware Information Systems for Emergency Management. In *ServiceWave Workshops*, pp. 50–58 (2010). doi: [10.1007/978-3-642-22760-8_5](https://doi.org/10.1007/978-3-642-22760-8_5).
- [32] DEARDEN, R., MEULEAU, N., RAMAKRISHNAN, S., SMITH, D., AND WASHINGTON, R. Incremental Contingency Planning. In *ICAPS'03 Workshop on Planning under Uncertainty and Incomplete Information*, pp. 38–47 (2003).
- [33] DI CICCIO, C., MARRELLA, A., AND RUSSO, A. Knowledge-intensive Processes: An Overview of Contemporary Approaches. In *1st International Workshop on Knowledge-intensive Business Processes (KiBP 2012)* (2012). Available from: http://ceur-ws.org/Vol-861/KiBP2012_paper_2.pdf.
- [34] DIJKMAN, R. M., DUMAS, M., AND OUYANG, C. Semantics and analysis of business process models in BPMN. *Information and Software Technology*, **50** (2008).
- [35] DUMAS, M. AND TER HOFSTEDÉ, A. H. M. UML Activity Diagrams as a Workflow Specification Language. In *Proceedings of the 4th International Conference on The Unified Modeling Language, Modeling Languages, Concepts, and Tools*, pp. 76–90 (2001).

- [36] DUMAS, M., VAN DER AALST, W., AND TER HOFSTEDÉ, A. *Process-aware information systems*. Wiley Online Library (2005).
- [37] EDELKAMP, S. AND HOFFMANN, J. PDDL2.2: The Language for the Classical Part of the 4th International Planning Competition. Tech. rep., Albert-Ludwigs-Universität Freiburg, Institut für Informatik (2004).
- [38] EDER, J. AND LIEBHART, W. The Workflow Activity Model WAMO. In *3rd international conference on Cooperative Information Systems (CoopIs)*, pp. 87–98 (1995).
- [39] EDER, J. AND LIEBHART, W. Workflow Recovery. In *Proceedings of the First IFCIS International Conference on Cooperative Information Systems, COOPIS '96*. IEEE Computer Society (1996).
- [40] FERREIRA, H. AND FERREIRA, D. An Integrated Life Cycle for Workflow Management Based on Learning and Planning. *Int. J. Cooperative Information Systems*, **15** (2006).
- [41] FOX, M. AND LONG, D. PDDL2.1: an Extension to PDDL for Expressing Temporal Planning Domains. *J. Artif. Int. Res.*, **20** (2003).
- [42] FRIEDRICH, G., FUGINI, M., MUSSI, E., PERNICI, B., AND TAGNI, G. Exception Handling for Repair in Service-Based Processes. *IEEE Transactions on Software Engineering*, **36** (2010).
- [43] GAJEWSKI, M., MEYER, H., MOMOTKO, M., SCHUSCHEL, H., AND WESKE, M. Dynamic Failure Recovery of Generated Workflows. In *Proceedings of the 16th International Workshop on Database and Expert Systems Applications (DEXA)*, pp. 982–986. IEEE Computer Society Press (2005).
- [44] GEREVINI, A., SAETTI, A., AND SERINA, I. Planning through stochastic local search and temporal action graphs in lpg. *J. Artif. Int. Res.*, **20** (2003).
- [45] GEREVINI, A., SAETTI, A., SERINA, I., AND TONINELLI, P. Lpg-td: a fully automated planner for pddl2.2 domains. In *In Proc. of the 14th Int. Conference on Automated Planning and Scheduling (ICAPS-04) International Planning Competition abstracts* (2004).
- [46] GHALLAB, M., AERONAUTIQUES, C., ISI, C. K., WILKINS, D., ET AL. Pddl—the planning domain definition language. Tech. rep. (1998).
- [47] GODEFROID, P. *Partial-Order Methods for the Verification of Concurrent Systems: An Approach to the State-Explosion Problem* (1996).
- [48] GREINER, U., RAMSCH, J., HELLER, B., LOFFLER, M., MULLER, R., AND RAHM, E. Adaptive guideline-based treatment workflows with adaptflow. *Studies in health technology and informatics*, (2004).
- [49] HALLERBACH, A., BAUER, T., AND REICHERT, M. Capturing variability in business process models: the provop approach. *Journal of Software Maintenance and Evolution: Research and Practice*, **22** (2009).

- [50] HALLERBACH, A., BAUER, T., AND REICHERT, M. Configuration and management of process variants. In *Handbook on Business Process Management 1* (edited by J. Brocke and M. Rosemann), International Handbooks on Information Systems. Springer Berlin Heidelberg (2010). ISBN 978-3-642-00415-5.
- [51] HAMADI, R., BENATALLAH, B., AND MEDJAHED, B. Self-adapting recovery nets for policy-driven exception handling in business processes. *Distributed and Parallel Databases*, **23** (2008).
- [52] HAMMER, M. *The Reengineering Revolution*. HarperCollins (1995). ISBN 9780887307362.
- [53] HELAL, S., MANN, W., EL-ZABADANI, H., KING, J., KADDOURA, Y., AND JANSEN, E. The gator tech smart house: A programmable pervasive space. *Computer*, **38** (2005).
- [54] HELMERT, M. New complexity results for classical planning benchmarks. In *Proceedings of the sixteenth international conference on automated planning and scheduling (ICAPS 2006)*, pp. 52–61 (2006).
- [55] HOLLINGSWORTH, D. ET AL. Workflow management coalition: The workflow reference model. *Document Number TC00-1003*, (1995).
- [56] HULL, R. Artifact-Centric Business Process Models: Brief Survey of Research Results and Challenges. In *Proceedings of the OTM 2008 Confederated International Conferences*, OTM '08, pp. 1152–1163 (2008).
- [57] HUMAYOUN, S. R., CATARCI, T., DE LEONI, M., MARRELLA, A., MECELLA, M., BORTENSLAGER, M., AND STEINMANN, R. Designing Mobile Systems in Highly Dynamic Scenarios. The WORKPAD Methodology. *Springer's International Journal on Knowledge, Technology and Policy*, **22** (2009). doi: [10.1007/s12130-009-9070-3](https://doi.org/10.1007/s12130-009-9070-3).
- [58] HUMAYOUN, S. R., CATARCI, T., DE LEONI, M., MARRELLA, A., MECELLA, M., BORTENSLAGER, M., AND STEINMANN, R. The WORKPAD User Interface and Methodology: Developing Smart and Effective Mobile Applications for Emergency Operators. In *5th International Conference on Universal Access in Human-Computer Interaction (UAHCI 2009)*, pp. 343–352. Springer Berlin Heidelberg (2009). doi: [10.1007/978-3-642-02713-0_36](https://doi.org/10.1007/978-3-642-02713-0_36).
- [59] I. BIDER. Masking Flexibility Behind Rigidity: Notes on How Much Flexibility People are Willing to Cope With. In *Proceedings of CAiSE05 Workshops*, pp. 7–18 (2005).
- [60] IBM INC. An introduction to WebSphere Process Server and WebSphere Integration Developer. Retrieved May 30, 2013, from: <ftp://ftp.software.ibm.com/software/integration/wps/library/WSW14021-USEN-01.pdf> (2008).
- [61] JARVIS, P., MOORE, J., STADER, J., MACINTOSH, A., DU MONT, A. C., AND CHUNG, P. Exploiting AI Technologies to Realise Adaptive Workflow

- Systems. *Proceedings of the AAAI Workshop on Agent-Based Systems in the Business Context*, (1999).
- [62] KALDELI, E., LAZOVIK, A., AND AIELLO, M. Continual planning with sensing for web service composition. In *AAAI* (2011).
- [63] KEMSLEY, S. The changing nature of work: From structured to unstructured, from controlled to social. In *BPM*, p. 2 (2011).
- [64] KLEIN, M., DELLAROCAS, C., AND BERNSTEIN, A. Introduction to the Special Issue on Adaptive Workflow Systems. *Computer Supported Cooperative Work (CSCW)*, **9** (2000).
- [65] KOWALSKI, R. A. ET AL. Using meta-logic to reconcile reactive with rational agents. *Meta-logics and logic programming*, (1995).
- [66] KÜNZLE, V. AND REICHERT, M. PHILharmonicFlows: towards a framework for object-aware process management. *Journal of Software Maintenance and Evolution: Research and Practice*, **23** (2011).
- [67] KÜNZLE, V. AND REICHERT, M. Striving for Object-aware Process Support: How Existing Approaches Fit Together. In *1st Int'l Symposium on Data-driven Process Discovery and Analysis (SIMPDA'11)* (2011).
- [68] LEAKE, D. B. *Case-based reasoning*. John Wiley and Sons Ltd. (2003).
- [69] LENZ, R. AND REICHERT, M. It support for healthcare processes - premises, challenges, perspectives. *Data Knowl. Eng.*, **61** (2007).
- [70] LEONE, N., PFEIFER, G., FABER, W., EITER, T., GOTTLOB, G., PERRI, S., AND SCARCELLO, F. The dl_v system for knowledge representation and reasoning. *ACM Transactions on Computational Logic (TOCL)*, **7** (2006).
- [71] LERNER, B., CHRISTOV, S., OSTERWEIL, L., BENDRAOU, R., KANNENGISSER, U., AND WISE, A. Exception handling patterns for process modeling. *Software Engineering, IEEE Transactions on*, **36** (2010).
- [72] LEVESQUE, H. J., REITER, R., LESPERANCE, Y., LIN, F., AND SCHERL, R. B. Golog: A logic programming language for dynamic domains. *The Journal of Logic Programming*, **31** (1997).
- [73] LEYMAN, F. AND ROLLER, D. *Production workflow: concepts and techniques*. Prentice Hall (2000).
- [74] LIU, C. AND CONRADI, R. Automatic replanning of task networks for process model evolution in EPOS. In *Software Engineering - ESEC'93*, pp. 434–450. Springer Berlin Heidelberg (1993).
- [75] LU, R. AND SADIQ, S. A survey of comparative business process modeling approaches. In *Business Information Systems*, Lecture Notes in Computer Science, pp. 82–94. Springer (2007). ISBN 978-3-540-72034-8.

- [76] MARRELLA, A. AND LESPERANCE, Y. Synthesizing a Library of Process Templates through Partial-Order Planning Algorithms. In *14th International Working Conference on Business Process Modeling, Development and Support (BPMDS 2013)*, pp. 277–291. Springer Berlin Heidelberg (2013). doi:10.1007/978-3-642-38484-4_20.
- [77] MARRELLA, A. AND MECELLA, M. Continuous Planning for Solving Business Process Adaptivity. In *12th International Working Conference on Business Process Modeling, Development and Support (BPMDS 2011)*, pp. 118–132 (2011). doi:10.1007/978-3-642-21759-3_9.
- [78] MARRELLA, A., MECELLA, M., AND RUSSO, A. Collaboration On-the-field: Suggestions and Beyond. In *8th International Conference on Information Systems for Crisis Response and Management (ISCRAM 2011)* (2011).
- [79] MARRELLA, A., MECELLA, M., AND RUSSO, A. Featuring Automatic Adaptivity through Workflow Enactment and Planning. In *7th International Conference on Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom 2011)*, pp. 372–381 (2011). doi:10.4108/icst.collaboratecom.2011.247096.
- [80] MARRELLA, A., MECELLA, M., RUSSO, A., TER HOFSTEDÉ, A. H. M., AND SARDIÑA, S. Making YAWL and SmartPM Interoperate: Managing Highly Dynamic Processes by Exploiting Automatic Adaptation Features. In *9th International Conference on Business Process Management (BPM 2011), Demonstration Track* (2011). Available from: <http://ceur-ws.org/Vol-820/Demo6.pdf>.
- [81] MARRELLA, A., RUSSO, A., AND MECELLA, M. Planlets: Automatically Recovering Dynamic Processes in YAWL. In *20th International Conference on Cooperative Information Systems (CoopIS 2012) - OTM Conferences (1)*, pp. 268–286 (2012). doi:10.1007/978-3-642-33606-5_17.
- [82] MCCARTHY, J. AND HAYES, P. J. Some philosophical problems from the standpoint of artificial intelligence. *Machine Intelligence*, **4** (1969).
- [83] MINOR, M., BERGMANN, R., AND GÖRG, S. Case-based adaptation of workflows. *Information Systems*, (2012).
- [84] MINOR, M., TARTAKOVSKI, A., AND SCHMALEN, D. Agile workflow technology and case-based change reuse for long-term processes. *International Journal of Intelligent Information Technologies (IJIIT)*, **4** (2008).
- [85] M.KINATEDER. Sap advanced workflow techniques. Retrieved May 30, 2013, from: <http://scn.sap.com/docs/DOC-3286>.
- [86] MONTALI, M., PESIC, M., VAN DER AALST, W. M. P., CHESANI, F., MELLO, P., AND STORARI, S. Declarative specification and verification of service choreographies. *ACM Transactions on the Web (TWEB)*, **4** (2010).

- [87] MÜLLER, R., GREINER, U., AND RAHM, E. AGENT WORK: a workflow system supporting rule-based workflow adaptation. *Data & Knowledge Engineering*, **51** (2004).
- [88] MURATA, T. Petri nets: Properties, analysis and applications. *Proceedings of the IEEE*, **77** (1989).
- [89] MYERS, K. AND BERRY, P. Workflow Management Systems: An AI Perspective. *AIC-SRI report*, (1998).
- [90] NAU, D., GHALLAB, M., AND TRAVERSO, P. *Automated Planning: Theory & Practice*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (2004).
- [91] NIGAM, A. AND CASWELL, N. S. Business Artifacts: An approach to Operational Specification. *IBM Systems Journal*, **42** (2003).
- [92] OASIS. Web Services Business Process Execution Language (WS-BPEL) ver. 2.0. <http://docs.oasis-open.org/wsbpel/2.0/wsbpel-v2.0.html> (2007).
- [93] OBERWEIS, A. Person-to-application processes: Workflow management. In *Process-Aware Information Systems* (2005).
- [94] OBJECT MANAGEMENT GROUP (OMG). Case Management Process Modeling (CMPM) – Request For Proposal. <http://www.omg.org/cgi-bin/doc?bmi/2009-9-23> (2009).
- [95] OUVANS, C., DUMAS, M., TER HOFSTEDÉ, A. H., AND VAN DER AALST, W. M. From bpmn process models to bpel web services. In *Web Services, 2006. ICWS'06. International Conference on*, pp. 285–292. IEEE (2006).
- [96] PESIC, M., SCHONENBERG, H., AND VAN DER AALST, W. M. P. DECLARE: Full Support for Loosely-Structured Processes. In *Proceedings of the 11th IEEE International Enterprise Distributed Object Computing Conference*, pp. 287–300 (2007).
- [97] PESIC, M. AND VAN DER AALST, W. M. P. A Declarative Approach for Flexible Business Processes Management. In *Business Process Management Workshops*, vol. 4103 of *LNCS*, pp. 169–180. Springer Berlin / Heidelberg (2006).
- [98] PETERSON, J. L. *Petri Net Theory and the Modeling of Systems*. Prentice Hall PTR, Upper Saddle River, NJ, USA (1981). ISBN 0136619835.
- [99] R-MORENO, M. D., BORRAJO, D., CESTA, A., AND ODDI, A. Integrating planning and scheduling in workflow domains. *Expert Systems with Applications: An International Journal*, **33** (2007).
- [100] R-MORENO, M. D. AND KEARNEY, P. Integrating AI planning techniques with Workflow Management System. *Knowl.-Based Syst.*, **15** (2002).
- [101] R-MORENO, M. D., ODDI, A., BORRAJO, D., AND CESTA, A. IPSS: A Hybrid Approach to Planning and Scheduling Integration. *Trans. on Kn. and Data Eng.*, **18** (2006).

- [102] REICHERT, M. What bpm technology can do for healthcare process support. In *AIME*, pp. 2–13 (2011).
- [103] REICHERT, M. AND DADAM, P. Adept_{flex}-supporting dynamic changes of workflows without losing control. *J. Intell. Inf. Syst.*, **10** (1998).
- [104] REICHERT, M., RINDERLE, S., KREHER, U., AND DADAM, P. Adaptive process management with adept2. In *ICDE*, pp. 1113–1114 (2005).
- [105] REICHERT, M. AND WEBER, B. *Enabling Flexibility in Process-Aware Information Systems - Challenges, Methods, Technologies*. Springer (2012). ISBN 978-3-642-30408-8.
- [106] REITER, R. Readings in Nonmonotonic Reasoning. chap. On Closed World Data Bases, pp. 300–310. Morgan Kaufmann Publishers Inc. (1987). ISBN 0-934613-45-1.
- [107] REITER, R. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. MIT Press (2001). ISBN 978-0262182188.
- [108] RINDERLE, S. AND REICHERT, M. Data Driven Process Control and Exception Handling in Process Management Systems. In *Advanced Information Systems Engineering* (edited by E. Dubois and K. Pohl), vol. 4001 of *LNCS*, pp. 273–287. Springer Berlin / Heidelberg (2006).
- [109] RINDERLE, S., WEBER, B., REICHERT, M., AND WILD, W. Integrating Process Learning and Process Evolution - A Semantics Based Approach. In *Business Process Management* (edited by W. van der Aalst, B. Benatallah, F. Casati, and F. Curbera), vol. 3649 of *LNCS*, pp. 252–267. Springer Berlin / Heidelberg (2005).
- [110] ROSEMAN, M. AND RECKER, J. Context-aware process design exploring the extrinsic drivers for process flexibility. In *BPMDs* (2006).
- [111] RUSSELL, N., TER HOFSTED, A., EDMOND, D., AND VAN DER AALST, W. M. P. Workflow Data Patterns: Identification, Representation and Tool Support. In *ER'05* (2005).
- [112] RUSSELL, N., VAN DER AALST, W., TER HOFSTED, A., AND EDMOND, D. Workflow resource patterns: Identification, representation and tool support. In *Advanced Information Systems Engineering*, pp. 11–42. Springer (2005).
- [113] RUSSELL, N., VAN DER AALST, W. M. P., AND TER HOFSTED, A. H. M. Workflow Exception Patterns. In *Advanced Information Systems Engineering* (edited by E. Dubois and K. Pohl), vol. 4001 of *LNCS*, pp. 288–302. Springer Berlin / Heidelberg (2006).
- [114] RUSSO, A., MECCELLA, M., AND DE LEONI, M. ROME4EU - A service-oriented process-aware information system for mobile devices. *Softw., Pract. Exper.*, **42** (2012).

- [115] SADIQ, S. W., SADIQ, W., AND ORLOWSKA, M. E. Pockets of flexibility in workflow specification. In *Proceedings of the 20th International Conference on Conceptual Modeling: Conceptual Modeling, ER '01*, pp. 513–526 (2001).
- [116] SCHMIDT, R. Flexible Support of Inter-Organizational Business Processes Using Web Services. In *Proceedings of CAiSE05 Workshops*, pp. 51–58 (2005).
- [117] SCHONENBERG, H., MANS, R., RUSSELL, N., MULAR, N., AND VAN DER AALST, W. M. P. Process flexibility: A survey of contemporary approaches. In *CIAO! / EOMAS* (2008).
- [118] SELMAN, B., KAUTZ, H. A., AND COHEN, B. Noise strategies for improving local search. In *Proceedings of the twelfth national conference on Artificial intelligence (vol. 1), AAAI '94*, pp. 337–343. American Association for Artificial Intelligence, Menlo Park, CA, USA (1994). ISBN 0-262-61102-3.
- [119] SHAZIA, W. S., MARJANOVIC, O., AND MARIA, E. O. Managing change and time in dynamic workflow processes. *International Journal of Cooperative Information Systems*, **9** (2000).
- [120] SMITH, D. E. AND WELD, D. S. Conformant graphplan. In *AAAI/IAAI*, pp. 889–896 (1998).
- [121] SOFFER, P. On the Notion of Flexibility in Business Processes. In *Proceedings of CAiSE05 Workshops*, pp. 26–37 (2005).
- [122] STAINES, T. S. Intuitive Mapping of UML 2 Activity Diagrams into Fundamental Modeling Concept Petri Net Diagrams and Colored Petri Nets. In *Proceedings of the 15th Annual IEEE International Conference and Workshop on the Engineering of Computer Based Systems*, pp. 191–200 (2008).
- [123] STRONG, D. M. AND MILLER, S. M. Exceptions and exception handling in computerized information processes. *ACM Trans. Inf. Syst.*, **13** (1995).
- [124] TER HOFSTEDÉ, A. H. M., VAN DER AALST, W. M. P., ADAMS, M., AND RUSSELL, N. *Modern Business Process Automation: YAWL and its Support Environment*. Springer (2009).
- [125] TIBCO SOFTWARE INC. TIBCO iProcess Engine - Architecture Guide. Retrieved May 30, 2013, from: <https://docs.tibco.com/pub/iprocess-engine/11.1.0-september-2009/pdf/tib-iprocess-engine-architecture-guide.pdf> (2009).
- [126] VAN BEEST, N., BULANOV, P., WORTMANN, H., AND LAZOVIK, A. Resolving business process interference via dynamic reconfiguration. In *8th International Conference on Service-Oriented Computing (ICSOC'10)*, pp. 47–60 (2010).
- [127] VAN BEEST, N., KALDELI, E., BULANOV, P., WORTMANN, H., AND LAZOVIK, A. Automatic detection of business process interference. In *1st International Workshop on Knowledge-intensive Business Processes (KiBP'12)*, pp. 6–20 (2012).

- [128] VAN DER AALST, W. The Application of Petri Nets to Workflow Management. *Journal of Circuits, Systems, and Computers*, (1998).
- [129] VAN DER AALST, W., ADAMS, M., TER HOFSTEDÉ, A., PESIC, M., AND SCHONENBERG, H. Flexibility as a service. In *Database Systems for Advanced Applications*, pp. 319–333. Springer (2009).
- [130] VAN DER AALST, W. AND VAN HEE, K. *Workflow Management: Models, Methods, and Systems*. MIT Press (2004).
- [131] VAN DER AALST, W. M. P. Formalization and verification of Event-driven Process Chains. *Information and Software Technology*, **41** (1999).
- [132] VAN DER AALST, W. M. P. Exterminating the dynamic change bug: A concrete approach to support workflow change. *Information Systems Frontiers*, **3** (2001).
- [133] VAN DER AALST, W. M. P., ALDRED, L., DUMAS, M., AND TER HOFSTEDÉ, A. H. M. Design and Implementation of the YAWL System. In *Proc. CAiSE'04*, vol. 3084 of *LNCS*, pp. 142–159 (2004).
- [134] VAN DER AALST, W. M. P. AND PESIC, M. DecSerFlow: Towards a Truly Declarative Service Flow Language. In *Web Services and Formal Methods*, vol. 4184 of *LNCS*, pp. 1–23. Springer Berlin / Heidelberg (2006).
- [135] VAN DER AALST, W. M. P., PESIC, M., AND SCHONENBERG, H. Declarative workflows: Balancing between flexibility and support. *Computer Science - R&D*, **23** (2009).
- [136] VAN DER AALST, W. M. P., TER HOFSTEDÉ, A., KIEPUSZEWSKI, B., AND BARROS, A. P. Workflow Patterns. *Distrib. Parallel Databases*, **14** (2003).
- [137] VARDI, M. An automata-theoretic approach to linear temporal logic. *Logics for concurrency*, (1996).
- [138] WEBER, B., REICHERT, M., AND RINDERLE-MA, S. Change Patterns and Change Support Features - Enhancing Flexibility in Process-aware Information Systems. *Data Knowl. Eng.*, **66** (2008).
- [139] WEBER, B., WILD, W., AND BREU, R. Cbrflow: Enabling adaptive workflow management through conversational case-based reasoning. *Advances in Case-Based Reasoning*, (2004).
- [140] WELD, D. S. An introduction to least commitment planning. *AI Magazine*, **15** (1994).
- [141] WESKE, M. Formal Foundation and Conceptual Design of Dynamic Adaptations in a Workflow Management System. In *Proceedings of the 34th Annual Hawaii International Conference on System Sciences (HICSS)* (2001).
- [142] WESKE, M. *Business Process Management: Concepts, Languages, Architectures*. Springer-Verlag New York, Inc., Secaucus, NJ, USA (2007).

-
- [143] WHITE, S. A. AND MIERS, D. *BPMN Modeling and Reference Guide: Understanding and Using BPMN*. Future Strategies Inc. (2008).
- [144] WIELEMAKER, J. An overview of the swi-prolog programming environment. In *WLPE*, pp. 1–16 (2003).
- [145] WILKINS, D. E., MYERS, K. L., LOWRANCE, J. D., AND WESLEY, L. P. Planning and reacting in uncertain and dynamic environments. *Journal of Experimental and Theoretical AI*, (1994).
- [146] WORKFLOW MANAGEMENT COALITION (WFMC). XPD L 2.1 Complete Specification. <http://www.wfmc.org/xpdl.html> (2008).
- [147] WORKFLOW MANAGEMENT COALITION (WFMC). Adaptive Case Management. <http://www.xpdl.org/nugen/p/adaptive-case-management/public.htm> (2010).
- [148] ZUR MUEHLEN, M. *Workflow-based process controlling: foundation, design, and application of workflow-driven process information systems*, vol. 6. Michael zur Muehlen (2004).