

Adaptive Process Management in Highly Dynamic and Pervasive Scenarios

Massimiliano de Leoni

Dipartimento di Informatica e Sistemistica

SAPIENZA - Università di Roma

deleoni@dis.uniroma1.it

Process Management Systems (PMSs) are currently more and more used as a supporting tool for cooperative processes in pervasive and highly dynamic situations, such as emergency situations, pervasive healthcare or domotics/home automation. But in all such situations, designed processes can be easily invalidated since the execution environment may change continuously due to frequent unforeseeable events. This paper aims at illustrating the theoretical framework and the concrete implementation of SmartPM, a PMS that features a set of sound and complete techniques to automatically cope with unplanned exceptions. PMS SmartPM is based on a general framework which adopts the Situation Calculus and IndiGolog.

1 Introduction

Nowadays organisations are always trying to improve the performance of the processes they are part of. It does not matter whether such organisations are dealing with classical static business domains, such as loans, bank accounts or insurances, or with pervasive and highly dynamic scenarios. The demands are always the same: seeking more efficiency for their processes to reduce the time and the cost for their execution.

According to the definition given by the Workflow Management Coalition¹, a workflow is “the computerised facilitation of automation of a business process, in whole or part”. The Workflow Management Coalition defines a Workflow Management System as “a system that completely defines, manages and executes workflows through the execution of software whose order of execution is driven by a computer representation of the workflow logic”. Workflow Management Systems (WfMSs) are also known as Process Management Systems (PMSs), and we are going to use both of them interchangeably throughout this thesis. Accordingly, this thesis uses many times word “process” in place of word “workflow”, although the original acceptance of the former is not intrinsically referring to its computerised automation.

In this paper we turn our attention to highly dynamic and pervasive scenarios. Pervasive scenarios comprise, for instance, emergency management, health care or home automation (a.k.a. domotics). All of these scenarios are characterised as being very dynamic and turbulent and subject to an higher frequency of unexpected contingencies with respect to classical scenarios. Therefore, PMSs for pervasive scenarios should provide a higher degree of operational flexibility/adaptability.

According to Andresen and Gronau [1] adaptability can be seen as an ability to change something to fit to occurring changes. Adaptability is to be understood here as the ability of a PMS to adapt/modify processes efficiently and fast to change circumstances. Adaptation aims at reducing the gap of the *virtual reality*, the (idealized) model of reality that is used by the PMS to deliberate, from the *physical reality*, the real world with the actual values of conditions and outcomes [2]. Exogenous events may make deviate

¹<http://wfmc.org>

Product	Manual	Pre-planned	Unplanned
YAWL		✓	
COSA	✓	✓	
Tibco	✓	✓	
WebSphere	✓	✓	
SAP	✓	✓	
OPERA	✓	✓	
ADEPT2	✓		
ADOME	✓		
AgentWork	✓		

Table 1: Adaptability in the leading PMSs (as from [11])

the virtual reality from the physical reality. The reduction of this gap requires sufficient knowledge of both kinds of realities (virtual and physical). Such knowledge, harvested by the services performing the process tasks, would allow the PMS to sense deviations and to deal with their mitigation.

In pervasive settings, efficiency and effectiveness when carrying on processes are a strong requirement. For instance, in emergency management saving minutes could result in saving injured people, preventing buildings from collapses, and so on. Or, pervasive health-care processes can cause people's permanent diseases when not executed by given deadlines. In order to improve effectiveness of process execution, adaptation ought to be as automatic as possible and to require minimum manual human intervention. Indeed, human intervention would cause delays, which might not be acceptable.

In theory there are three possibilities to deal with deviations:

1. Ignoring deviations – this is, of course, not feasible in general, since the new situation might be such that the PMS is no more able to carry out the process instance.
2. Anticipating all possible discrepancies – the idea is to include in the process schema the actions to cope with each of such failures. This can be seen as a *try-catch* approach, used in some programming languages such as Java. The process is defined as if exogenous actions cannot occur, that is everything runs fine (the *try* block). Then, for each possible exogenous event, a *catch* block is designed in which the method is given to handle the corresponding exogenous event. For simple and mainly static processes, this is feasible and valuable; but, especially in mobile and highly dynamic scenarios, it is quite impossible to take into account all exception cases.
3. Devising a general recovery method able to handle any kind of exogenous events – considering again the metaphor of *try/catch*, there exists just one *catch* block, able to handle any exogenous events, included the unexpected. The *catch* block activates the general recovery method to modify the old process P in a process P' so that P' can terminate in the new environment and its goals are included in those of P . This approach relies on the execution monitor (i.e., the module intended for execution monitoring) that detects discrepancies leading the process instance not to be terminable. When they are sensed, the control flow moves to the *catch* block. An important challenge here is to build the monitor which is able to identify which exogenous events are relevant, i.e. that make impossible process to terminate, as well as to *automatically* synthesize P' during the execution itself.

Table 1 shows the adaptability features of the most valuable PMSs according to the state-of-art analysis described in [11]. Column **Manual** refers to the possibility of a responsible person who manually changes the process schema to deal with exogenous events. Column **Pre-planned** concerns the feature

of defining policies to specify the adaptation behaviour to manage some exogenous events, whose possible occurrence is foreseeable a priori. The last column **Unplanned** refers to the third approach in the classification above.

The third approach seems to be the most appropriate when dealing with scenarios where (i) the frequency of unexpected exogenous events are relatively high and (ii) there are several exogenous events that cannot be foreseen before their actual occurrence. Unfortunately, as the table shows, the world leading PMSs are unable to feature the third approach.

This paper describes SmartPM, a PMS that features some sound and complete techniques according to the third approach described above. Such techniques are meant to improve the degree of *automatic* adaptation to react to very frequent changes in the execution environment and fit processes accordingly. The techniques proposed here are based on Situation Calculus [13] and automatic planning, conceived to coordinate robots and intelligent agents. The concrete implementation, namely SmartPM, is based on the IndiGolog interpreter developed at University of Toronto and RMIT University, Melbourne.

In SmartPM, every entity performing task is generally named “service”. A service may be a human actor/process participant as well as an automatic service that execute a certain job (e.g., a SOAP-based Web Service).

Let us consider a scenario for emergency management where processes show typical a complexity that is comparable to business settings. Therefore, the usage of PMS is valuable to coordinate the activities of emergency operators. In these scenarios, operators are typically equipped with low-profile devices, such as PDAs, which several services are installed on. Such services may range from usual GUI-based applications to automatic ones. For instances, some applications can be installed to fill questionnaires or take pictures. In addition, PDAs can be provided with some automatic services that connect to the Civil Protection headquarters to retrieve information for the assessment of the affected area and possibly send back the data collected.

PDAs communicate with each other by Mobile Ad-hoc Networks (MANETs), which are Wi-Fi networks that do not rely on a fixed infrastructure, such as Access Points. Devices can be the final recipients of some packets sent by other devices as well as they can act as relays and forward packets towards the final destination.

In order to orchestrate the services installed on operator devices, such devices need to be continually connected to the PMS through a loose connection: devices and the PMS can communicate if there exists a path of nodes that connects them in the graph of the communication links.

In the virtual reality, devices are supposed to be continuously connected (i.e., a path always exists between pairs of nodes). But in this physical reality continuous connections cannot be guaranteed: the environment is highly dynamic and the movement of nodes (that is, devices and related operators) within the affected area, while carrying out assigned tasks, can cause disconnections and make deviate the two reality. Disconnections results in the unavailability of nodes and, hence, the services provided. From the collection of actual user requirements [6], it results that typical teams are formed by a few nodes (less than 10 units), and therefore frequently a simple task reassignment is not feasible. Indeed, there may not be two “similar” services available to perform a given task. Reordering task executions would not solve the problem, either. There is no guarantee that eventually those services that provide unique capability connect again to the PMS.

So, adaptation is needed: adaptability might consist in this case to recover the disconnection of a node X, and that can be achieved by assigning a task “Follow X” to another node Y in order to maintain the connection. When the connection has been restored, the process can progress again.

Construct	Meaning	Platform Statement
a	A primitive action	a
$\phi?$	Wait while the ϕ condition is false	$?(\phi)$
$(\delta_1; \delta_2)$	Sequence of two sub-programs δ_1 and δ_2	$[\delta_1, \delta_2]$
$proc P(\vec{v}) \delta$	Invocation of a procedure passing a vector \vec{v} of parameters	$proc(P, \delta)$
$(\phi; \delta_1) (\neg\phi; \delta_2)$	Exclusive choice between δ_1 and δ_2 according to the condition ϕ	$ndet([\phi; \delta_1], [\neg(\phi); \delta_2])$
$while \phi do \delta$	Iterative invocation of δ	$while(\phi, \delta)$
$(\delta_1 \delta_2)$	Concurrent execution	$rrobin(\delta_1, \delta_2)$
δ^*	Indeterministic iteration of program execution (The platform statement limits the maximum iterations number to n)	$star(\delta, n)$
$\Sigma(\delta)$	Emulating off-line execution	$searchn(\delta, n)$
$\pi a. \delta$	Indeterministic choice of argument a followed by the execution of δ	$pi(a, \delta)$

Table 2: IndiGolog constructs

2 Preliminaries

In this section we introduce the Situation Calculus, which we use to formalize SmartPM and its adaptation features. The Situation Calculus [13] is a second-order logic targeted specifically for representing a dynamically changing domain of interest (the world). All changes in the world are obtained as result of *actions*. A possible history of the actions is represented by a *situation*, which is a first-order term denoting the current situation of the world. The constant s_0 denotes the initial situation. A special binary function symbol $do(\alpha, s)$ denotes the next situation after performing the action α in the situation s . Action may be parameterized.

Properties that hold in a situation are called *fluents*. These are predicates taking a situation term as their last argument. For instance, we could define the fluent $free(x, s)$ stating whether the object x is free in situation s , meaning no object is located on x in situation s .

Changes in fluents (resulting from executing actions) are specified through *successor state axioms*. In particular for each fluent F we have a successor state axioms as follows:

$$F(\vec{x}, do(\alpha, s)) \Leftrightarrow \Phi_F(\vec{x}, do(\alpha, s), s)$$

where $\Phi_F(\vec{x}, do(\alpha, s), s)$ is a formula with free variables \vec{x} , α is an action, and s is a situation.

In order to control the executions of actions we make use of high level programs expressed in IndiGolog [14], which is equipped with primitives for expressing concurrency. Table 2 summarizes the constructs of IndiGolog used in this work. Basically, these constructs allow to define every well-structured process as defined in [7]. The last table column shows the corresponding statement defined in the IndiGolog platform developed at University of Toronto and RMIT University.²

From the formal point of view, IndiGolog programs are terms. The execution of ConGolog programs is expressed through a *transition semantic* based on single steps of execution. At each step a program executes an action and evolves to a new program which represents what remains to be executed of the original program. Formally two predicates are introduced to specify such a semantic:

²Downloadable at <http://www.cs.toronto.edu/cogrobo/main/systems/index.html>

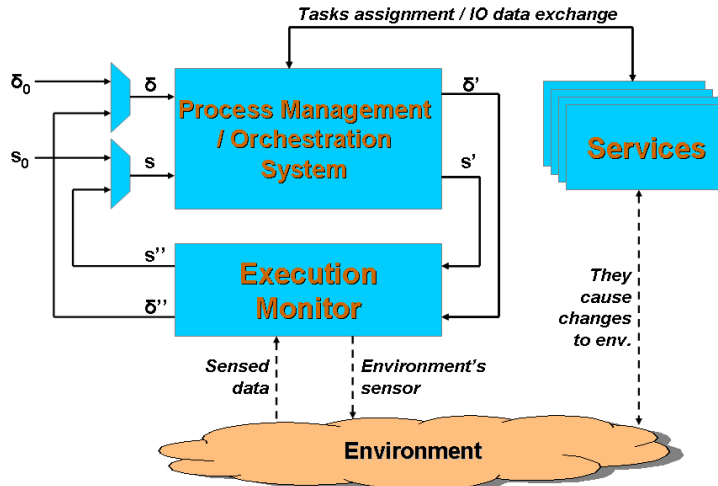


Figure 1: Execution Monitoring.

- $Trans(\delta', s', \delta'', s'')$, 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' .

By using $Trans$ and $Final$ we can define a predicate $Do(\delta', s', s'')$ that represent successful complete executions of a program δ' in a situation s' , where s'' is the situation at the end of the execution of δ' . Formally:

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

where $Trans^*$ is the definition of the reflective and transitive closure of $Trans$.

To cope with the impossibility of backtracking actions executed in the real world, IndiGolog incorporates a new programming construct, namely the *search operator*. 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 and possible picks of variable values, it performs reasoning in order to decide for a step which still allows the rest of δ to terminate successfully. If δ is the entire program under consideration, $\Sigma(\delta)$ emulates complete off-line execution.

3 General Framework

The general framework which we shall introduce in this paper is based on the *execution monitoring* scheme as described in [2] for situation calculus agents. As we will later describe in more details, when using IndiGolog for process management, we take tasks to be predefined sequences of actions (see later) and processes to be IndiGolog programs. After each action, the PMS may need to align the internal world representation (i.e., the virtual reality) with the external one (i.e., the physical reality).

Before a process starts, PMS takes the initial context from the real environment and builds the corresponding initial situation S_0 , by means of first-order logic formulas. It also builds the program δ_0 corresponding to the process to be carried on. Then, at each execution step, PMS, which has a complete knowledge of the internal world (i.e., its virtual reality), assigns a task to a service. The only “assignable”

tasks are those whose preconditions are fulfilled. A service can collect data required needed to execute the task assigned from PMS. When a service finishes executing a task, it alerts PMS of that.

The execution of the PMS can be interrupted by the *monitor* module when a misalignment between the virtual and the physical realities is discovered. In that case, the monitor *adapts* the (current) program to deal with such discrepancy.

In Figure 1, the overall framework is depicted. At each step, the PMS advances the process δ in situation s by executing an action, resulting then in a new situation s' with the process δ' remaining to be executed. Both s' and δ' are given as input to the monitor, which also collects data from the environment through *sensors*.³ If a discrepancy between the virtual reality as represented by s' and the physical reality is sensed, then the monitor changes s' to s'' , by generating a sequence of actions that explains the changes perceived in the environment, thus re-aligning the virtual and physical realities. Notice, however, that the process δ' may *fail* to execute successfully (i.e., assign all tasks as required) in the new (unexpected) situation s'' . If so, the monitor adapts also the (current) process by performing suitable recovery changes and generating then a new process δ'' . At this point, the PMS is resumed and the execution continues with program-process δ'' in situation s'' .

4 Process Formalisation in Situation Calculus

Next we detail the general framework proposed above by using Situation Calculus and IndiGolog. We use some domain-independent predicates to denote the various objects of interest in the framework:

- *service*(a): a is a service
- *task*(x): x is a task
- *capability*(b): b is a capability
- *provide*(a, b): the service a provides the capability b
- *require*(x, b): the task x requires the capability b

In the light of these predicates, we have defined a shortcut to refer to the capability of a certain service a to perform a list of tasks, a.k.a. worklist. Service a can execute a certain worklist *wrkList* iff a provides all capabilities required by all tasks in the worklist:

$$Capable(a, wrklist) \Leftrightarrow (\forall b, t. t \in wrkList \wedge require(b, t) \Rightarrow provide(a, b))$$

Every task execution is the sequence of four PMS actions: (i) the assignment of the task to a service, resulting in the service being not free anymore; (ii) the notification to the service to start executing the task. Then, the service carries out the tasks and, after receiving the service notification of the task conclusion, (iii) the PMS acknowledges the successful task termination. Finally, (iv) the PMS releases the service, which becomes free again. We formalise these four actions as follows:

- *Assign*(a, x): task x is assigned to a service a
- *Start*(a, x, p): service a is allowed to start the execution of task x . The input provided is p .
- *AckTaskCompletion*(a, x): service a concluded successfully the executing of x .

³Here, we refer as *sensors* not only proper sensors (e.g., the ones deployed in sensor networks), but also any software or hardware component enabling to retrieve contextual information. For instance, it may range from GIS clients to specific hardware that makes available the communication distance of a device to its neighbors. [10]

- $Release(a,x)$: the service a is released with respect to task x .

In addition, services can execute two actions:

- $readyToStart(a,x)$: service a declares to be ready to start performing task x
- $finishedTask(a,x,q)$: services declares to have completed the execution of task x returning output q .

The terms p and q denote arbitrary sets of input/output, which depend on the specific task. Special constant \emptyset denotes empty input or output.

The interleaving of actions performed by the PMS and services is as follows. After the assignment of a certain task x by $Assign(a,x)$, when the service a is ready to start executing, it executes action $readyToStartTask(a,x)$. At this stage, PMS executes action $Start(a,x,p)$, after which a starts executing task x . When a completes task x , it executes the action $finishedTask(a,x,q)$. Specifically, we envision that actions $finishedTask(\cdot)$ are those in charge of changing properties of world as result of executing tasks. When x is completed, PMS is allowed in any moment to execute sequentially $AckTaskCompletion(a,x)$ and $Release(a,x)$. The program coding the process will be executed by only one actor, specifically the PMS. Therefore, actions $readyToStartTask(\cdot)$ and $finishedTask(\cdot)$ are considered as external and, hence, not coded in the program itself.

For each specific domain, we have several fluents representing the properties of situations. Some of them are modelled independently of the domain whereas others, the majority, are defined according to the domain. If they are independent of the domain, they can be always formulated as defined in this chapter. Among the domain-independent ones, we have fluent $free(a,s)$, that denotes the fact that the service a is free, i.e., no task has been assigned to it, in the situation s . The corresponding successor state axiom is as follows:

$$\begin{aligned} free(a,do(t,s)) \Leftrightarrow & \\ & (\forall x.t \neq Assign(a,x) \wedge free(a,s)) \vee \\ & (\neg free(a,s) \wedge \exists x.t = Release(a,x)) \end{aligned} \quad (1)$$

This says that a service a is considered free in the current situation if and only if a was free in the previous situation and no tasks have been just assigned to it, or a was not free and it has been just released. There exists also the domain-independent fluent $enabled(x,a,s)$ which aims at representing whether service a has notified to be ready to execute a certain task x so as to enabled it. The corresponding successor-state axiom:

$$\begin{aligned} enabled(x,a,do(t,s)) \Leftrightarrow & \\ & (enabled(x,a,s) \wedge \forall q.t \neq finishedTask(a,x,q)) \vee \\ & (\neg enabled(x,a,s) \wedge t = readyToStartTask(a,x)) \end{aligned} \quad (2)$$

This says that $enabled(x,a,s)$ holds in the current situation if and only if it held in the previous one and no action $finishedTask(a,x,q)$ has been performed or it was false in the previous situation and $readyToStartTask(a,x)$ has been executed. This fluent aims at enforcing the constraints that the PMS can execute $Start(a,x,p)$ only after a performed $begun(a,x)$ and it can execute $AckTaskCompletion(a,x,q)$ only after $finishedTask(a,x,q)$. This can be represented by two pre-conditions on actions $Start(\cdot)$ and $AckTaskCompletion(\cdot)$:

$$\begin{aligned} \forall p.Poss(Start(a,x,p),s) \Leftrightarrow enabled(x,a,s) & \\ \forall p.Poss(AckTaskCompletion(x,a),s) \Leftrightarrow \neg enabled(x,a,s) & \end{aligned} \quad (3)$$

provided that $AckTaskCompletion(x,a)$ never comes before $Start(x,a,p),s$.

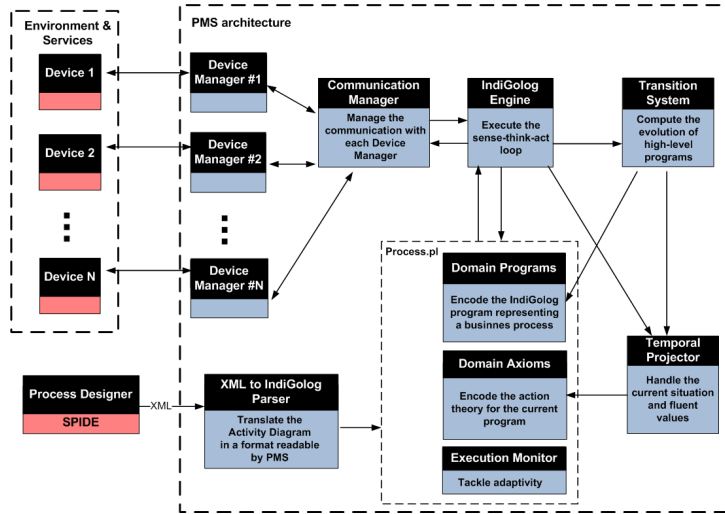


Figure 2: Architecture of the PMS.

Furthermore, we introduce a domain-independent fluent $started(x, a, p, s)$ that holds if and only if an action $Start(a, x, p)$ has been executed but the dual $AckTaskCompletion(x, a)$ has not yet:

$$\begin{aligned}
 started(a, x, p, do(t, s)) &\Leftrightarrow \\
 &(started(a, x, p, s) \wedge t \neq Stop(a, x)) \vee \\
 &(\nexists p'. started(x, a, p', s) \wedge t = Start(a, x, p))
 \end{aligned} \tag{4}$$

In addition, we make use, in every specific domain, of a predicate $available(a, s)$ which denotes whether a service a is available in situation s for tasks assignment. However, $available$ is domain-dependent and, hence, requires to be defined specifically for every domain. Knowing whether a service is available is very important for the PMS when it has to perform assignments. Indeed, a task x is assigned to the best service a which is available and provides every capability required by x . The fact that a certain service a is free does not imply it can be assigned to tasks (e.g., in the example described above it has to be free as well as it has to be indirectly connected to the coordinator). The definition of $available(\cdot)$ must enforce the following condition:

$$\forall a. s. available(a, s) \Rightarrow free(a, s) \tag{5}$$

We do not give explicitly pre-conditions to task. We assume tasks can always be executed. We assume that, given a task, if some conditions do not hold, then the outcomes of that tasks are not as expected (in other terms, it fails).

5 The SmartPM System

This section aims at describing the internal structure of PMS. Figure 2 shows its conceptual architecture. At the beginning, a responsible person designs an Activity Diagram through SPIDE, a *Process Designer* Graphical tool with which SmartPM is equipped. Later, Such a tool translates the Activity Diagram in a XML format file. Then, such a XML file is loaded into PMS. The *XML-to-IndiGolog Parser* component translates this specification in a *Domain Program*, the IndiGolog program corresponding to the designed

process, and a set of *Domain Axioms*, which is the action theory that comprises the initial situation, the set of available actions with their pre- and post-conditions.

When the program is translated in the Domain Program and Axioms, a component named *Communication Manager* (CM) starts up all of *device managers*, which are basically some drivers for making communicate PMS with the services and sensors installed on devices. For each real world device PMS holds a device manager. Each device manager is also intended for notifying the associated device about every action performed by the SmartPM engine as well as for notifying the SmartPM engine about the actions executed by the services of the associated device.

After this initialization process, CM activates the *IndiGolog Engine*, which is in charge of executing IndiGolog programs. Then, CM enters into a passive mode where it is listening for messages arriving from the devices through the device managers. In general, a message can be a exogenous event harvested by a certain sensor installed on a given device as well as a message notifying the start or completion of a certain task. When CM judges a message as significant, it forwards it to IndiGolog. For instance, relevant messages may be signals of the task completion or the sudden unavailability of a given device.

In sum, CM is responsible of deciding which device should perform certain actions, instructing the appropriate device managers to communicate with the device services and collecting the corresponding sensing outcome. The IndiGolog Engine is intended to execute a *sense-think-act* interleaved loop [8]. The cycle repeats at all times the following three steps:

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, instructing the Communication Manager.

The IndiGolog Engine relies on two further modules named *Transition System* and *Temporal Projector*. The former is used to compute the evolution of IndiGolog programs according to the statements' semantic, whereas the latter is in charge of holding the current situations throughout the execution as well as letting evaluate the fluent values for taking the right decision of the actions to perform.

The last module that is worth mentioning is the *Execution Monitor* (MON), which get notifications of exogenous events from the Communication Manager. It decides whether adaptation is needed and adapts accordingly the process. Section 7.2 gives some additional details of the concrete implementation of monitoring and adaptation.

6 A concrete example from Emergency Management

We turn to describe the approach by an example concerning emergency management in an area affected by an earthquake. The emergency response process in question comprises various activities that may need to be adapted on-the-fly to react to unexpected exogenous events that could arise during the operation. Figure 3 depicts an Activity Diagram of a process consisting of two concurrent branches; the final task is *send data* and can only be executed after the branches have successfully completed. The left branch, abstracted out from the diagram, is built from several concurrent processes involving tasks *rescue*, *evacuation* and others. The right branch begins with the concurrent execution of three sequences of tasks: *go*, *photo*, and *survey*. When all survey tasks have been completed, the task *evaluate pictures* is executed. Then, a condition is evaluated on the resulting state at a decision point (i.e., whether the pictures taken are of sufficient quality). If the condition holds, the right branch is considered finished; otherwise, the whole branch should be repeated.




Figure 3: An activity diagram of a process concerning emergency management.

Figure 4 shows some parts of the IndiGolog program representing the process of the example. The code proposed here has been slightly simplified and abstracted for the sake of brevity. The main procedure, called `main`, involves three interrupts running at different priorities. The first highest priority interrupt fires when an exogenous event occurs (i.e., condition `exogEvent` is true). In such a case, the `monitor` procedure is executed, evaluating whether or not adaptation is required (see Section 7.2).

If no exogenous event has occurred, the second interrupt triggers and execution of the actual emergency response process is attempted. Procedure `process`, also shown in the figure, encodes the Activity Diagram of the example process. It relies, in turn, on procedure `manageTasks(WrkLists)`, where `WrkLists` is a sequence of elements `workitem(T, I, D)`, each one representing a task `T`, with identifier `I`, and input data `D`, which needs to be performed. This procedure is meant to manage the execution of all tasks in the worklist, and it assigns them all to a *single* service that provides every capability required.

Of course, to assign tasks to an service, SmartPM needs to reason about the available ones, their current state (e.g., their location), and their capabilities, as not every service is capable of performing any task. In fact, before assigning the first task in any task list, procedure `manageTasks(WrkLists)` executes a *pick* operation to choose a Service `srvc` that is involved in no task execution (i.e.,

```

proc(main,
  prioritized_interrupts(
    [interrupt(exogEvent, monitor),
     interrupt(true, process),
     interrupt(neg(finished), wait)]
  )).
proc(process, [rrobin(processRescue,
  while(or(noPhotos<7,neg(goodPics)),
    [rrobin(
      [manageTasks(
        [workitem((go,id19,loc(5,5)),
         workitem((photo,id20,loc(5,5)),
          workitem((survey,id21,loc(5,5))]),
        manageTasks(
          [workitem((go,id19,loc(15,15)),
           workitem((photo,id20,loc(15,15)),
            workitem((survey,id21,loc(15,15))]),
          manageTasks(
            [workitem((go,id19,loc(50,50)),
             workitem((photo,id20,loc(50,50)),
              workitem((survey,id21,loc(50,50))]),
            ]
          ),
          manageTasks([workitem((evalPics,id28,input))]
        )] % end of while
      ), % end concurrent subprocesses
      manageTasks([workitem((sendData,id29,input))]
    )].

proc(manageTasks(WrkList),
  pi(srvc,
    [?(and(Available(srvc),Capable(srvc,WrkList))),
     manageExecution(WrkList,srvc),
    ]
  )).
proc(manageExecution([],Srvc,[]).
proc(manageExecution([workitem(Task,Id,I)|TAIL],Srvc),
  [assign(Task,Id,Srvc,I),
   start(Task,Id,Srvc,I),
   ackTaskCompletion(Task,Id,Srvc),
   release(Task,Id,Srvc,I),
   manageExecution(TAIL,Srvc)
  ]
)
)

```

Figure 4: An example of process management with IndiGolog.

fluent `Free(ctr)` holds) and able to execute the whole worklist.

Once a suitable service has been chosen, PMS assigns the list of tasks to it by executing `assign(srvc,WrkList)`. In addition to inform the service about the task assignment, such an action turns fluent `Free(ctr)` to false.

Then, PMS calls procedure `manageExecution(WrkList)`, which handles the execution of each task in the list. For each task `T` in the list (with identifier `I` and input data `D`), the procedure invokes action `start(T,D,I,srvc)` that provides the required information to the chosen service `srvc`. In this way, the service is instructed to begin working on the task and receives the required input. When a service finishes executing an assigned task, it alerts SmartPM via action `finishedTask(T,srvc)`; PMS acknowledges by performing `ackTaskCompletion(T,D,ctr)`. When the whole work-item list is execution, the PMS releases the service by executing the action `release(T,D,ctr)`, after which fluent `Free(srvc)` is turned to true again.

It is worth mentioning that, if the process being carried out cannot execute temporarily further, the lowest priority interrupt fires. This interrupt makes PMS wait for the conditions in which some tasks can be executed. The fact that the process gets stuck does not imply necessarily the occurrence of some relevant exogenous events. It could be also caused by the fact that next tasks can be only assigned to services that are currently busy busy performing other tasks. The latter situation does not prevent processes from being completed successfully; indeed, such services will be eventually free to work on those tasks.

7 Adaptation in SmartPM

7.1 Monitoring Formalisation

Next we formalize how the monitor works. Intuitively, the monitor takes the current program δ' and the current situation s' from the PMS's virtual reality and, analyzing the physical reality by sensors, introduces fake actions in order to get a new situation s'' which aligns the virtual reality of the PMS with sensed information. Then, it analyzes whether δ' can still be executed in s'' , and if not, it adapts δ' by generating a new correctly executable program δ'' . Specifically, the monitor work can be abstractly defined as follows (we do not model how the situation s'' is generated from the sensed information):

$$\begin{aligned} \text{Monitor}(\delta', s', s'', \delta'') \Leftrightarrow & (\text{Relevant}(\delta', s', s'') \wedge \text{Recovery}(\delta', s', s'', \delta'')) \vee \\ & (\neg \text{Relevant}(\delta', s', s'') \wedge \delta'' = \delta') \end{aligned} \quad (6)$$

where: (i) $\text{Relevant}(\delta', s', s'')$ states whether the change from the situation s' into s'' is such that δ' cannot be correctly executed anymore; and (ii) $\text{Recovery}(\delta', s', s'', \delta'')$ is intended to hold whenever the program δ' , to be originally executed in the situation s' , is adapted to δ'' in order to be executed in the situation s'' .

Formally Relevant is defined as follows:

$$\text{Relevant}(\delta', s', s'') \Leftrightarrow \neg \text{SameConfig}(\delta', s', \delta', s'')$$

where $\text{SameConfig}(\delta', s', \delta'', s'')$ is true if executing δ' in s' is “equivalent” to executing δ'' in s'' (see later for further details).

In this general framework we do not give a definition for $\text{SameConfig}(\delta', s', \delta'', s'')$. However we consider any definition for SameConfig to be correct if it denotes a bisimulation [12]. Formally, for every $\delta', s', \delta'', s''$ holds:

1. $\text{Final}(\delta', s') \Leftrightarrow \text{Final}(\delta'', s')$
2. $\forall a, \delta'. \text{Trans}(\delta', s', \overline{\delta'}, do(a, s')) \Rightarrow$
 $\exists \overline{\delta''}. \text{Trans}(\delta'', s'', \overline{\delta'}, do(a, s'')) \wedge \text{SameConfig}(\overline{\delta'}, do(a, s), \overline{\delta''}, do(a, s''))$
3. $\forall a, \delta'. \text{Trans}(\delta'', s'', \overline{\delta'}, do(a, s'')) \Rightarrow$
 $\exists \overline{\delta'}. \text{Trans}(\delta', s', \overline{\delta'}, do(a, s')) \wedge \text{SameConfig}(\overline{\delta''}, do(a, s''), \overline{\delta'}, do(a, s'))$

Intuitively, a predicate $\text{SameConfig}(\delta', s', \delta'', s'')$ is said to be correct if δ' and δ'' are terminable either both or none of them. Furthermore, for each action a performable by δ' in the situation s' , δ'' in the situation s'' has to enable the performance of the same actions (and viceversa). Moreover, the resulting configurations $(\overline{\delta'}, do(a, s'))$ and $(\overline{\delta''}, do(a, s''))$ must still satisfy SameConfig .

The use of the bisimulation criteria to state when a predicate $SameConfig(\dots)$ is correct, derives from the notion of equivalence introduced in [5]. When comparing the execution of two formally different business processes, the internal states of the processes may be ignored, because what really matters is the process behavior that can be observed. This view reflects the way a PMS works: indeed what is of interest is the set of tasks that the PMS offers to its environment, in response to the inputs that the environment provides.

Next we turn our attention to the procedure to adapt the process formalized by $Recovery(\delta, s, s', \delta')$. Formally is defined as follows:

$$Recovery(\delta', s', s'', \delta'') \Leftrightarrow \exists \delta_a, \delta_b. \delta'' = \delta_a; \delta_b \wedge Deterministic(\delta_a) \wedge Do(\delta_a, s'', s_b) \wedge SameConfig(\delta', s', \delta_b, s_b) \quad (7)$$

$Recovery$ determines a process δ'' consisting of a *deterministic* δ_a (i.e., a program not using the concurrency construct), and an arbitrary program δ_b . The aim of δ_a is to lead from the situation s'' in which adaptation is needed to a new situation s_b where $SameConfig(\delta', s', \delta_b, s_b)$ is true.

The nice feature of RECOVERY is that it asks to search for a linear program that achieves a certain formula, namely $SameState(s', s'')$. That is we have reduced the synthesis of a recovery program to a classical Planning problem in AI [4]. As a result we can adopt a well-developed literature about planning for our aim. In particular, if the services and input and output parameters are finite, then the recovery can be reduced to *propositional* planning, which is known to be decidable in general (for which very well performing software tools exists).

Notice that during the actual recovery phase δ_a we disallow for concurrency because we need full control on the execution of each service in order to get to a recovered state. Then the actual recovered program δ_b can again allow for concurrency.

In the previous sections we have provided a general description on how adaptation can be defined and performed. Here we choose a specific technique that is actually feasible in practice. Our main step is to adopt a specific definition for $SameConfig$, here denoted as SAMECONFIG, namely:

$$SAMECONFIG(\delta', s', \delta'', s'') \Leftrightarrow SameState(s', s'') \wedge \delta' = \delta'' \quad (8)$$

In other words, SAMECONFIG states that δ', s' and δ'', s'' are the same configuration if (i) all fluents have the same truth values in both s' and s'' ($SameState$), and (ii) δ'' is actually δ' .⁴ In papers [11, 9], we have proved that the above-defined SAMECONFIG is a correct bisimulation.

Using Equation 8 as $SameConfig$ definition feasible in practice, relevancy results to be:

$$RELEVANT(\delta', s', s'') \Leftrightarrow \neg SameState(s', s'') \quad (9)$$

In the next section, we are going to show how the abstract planner specification given here has been concretely used inside SmartPM. Specifically the current version of SmartPM uses the proportional planner available in the IndiGolog platform developed by University of Toronto and RMIT in Melbourne. In order to adapt, SmartPM is based on the concrete definitions of relevancy and $SameConfig$ given by Equations 9 and 8.

7.2 The execution monitoring and adaptation

⁴Observe that $SameState$ can actually be defined as a first-order formula over the fluents, as the conjunction of $F(s') \Leftrightarrow F(s'')$ for each fluent F .

```

proc(monitor, [ndet(
    [?(neg(relevant))],
    [?(relevant),recovery]
)]).

proc(recovery, searchn([searchProgram],10).

proc(searchProgram, [star(pi([Task,Id,Input,srv],
    [?(and(Available(srv),
        Capable(srv,[workitem(Task,Id,Input]))]),
        manageExecution([workitem(Task,Id,Input],srv))]),
    ?(SameState))]).

```

Figure 5: The procedure for managing automatic adaptation with the IndiGolog interpreter.

As already told, adaptation amounts to find a linear program (i.e., without concurrency) that is meant to be “appended” before the current IndiGolog program remaining to be executed. Such a linear program is meant to resolve the gap that was just sensed by restoring the values of affected fluents to those before the occurrence of the deviation.

Figure 5 shows how adaptability has been concretely implemented in SmartPM. The execution of the process being carried out by SmartPM can be interrupted by the `monitor` procedure when a misalignment between the virtual and the physical reality is discovered.

The `monitor` procedure is the concrete coding of Equation 6 and relies on procedure `relevant`. Procedure `relevant` returns true if the exogenous event has created a gap between the physical and virtual reality that is in accord with Equation 9. For this aim, SmartPM keeps a “copy” of the expected value of each defined fluent so that when an exogenous action is sensed it can check whether the action has altered the value of some fluent.

If the gap is relevant, procedure `recovery` is invoked. It amounts to find a linear program (i.e., without concurrency) to reduce the gap sensed as well as, if such a program is found, to execute it. After executing such a linear program, the program coded by routine `process` (and its possible sub-routines) can progress again. This behaviour is equivalent to that expressed formally in Equation 7 where the adapting linear program is “appended before” and, hence, executed before the remaining process.

The `recovery` procedure looks for a sequence of actions that brings to a situation in which procedure `SameState` returns true: $\Sigma((\pi a.a)^*; SameState?)$. Procedure `SameState` tests whether executing $(\pi a.a)^*$ really has really reduced the gap. The use of the IndiGolog’s lookahead operator Σ guarantees the action sequence $(\pi a.a)^*$ is chosen so as to make `SameState` true. In fact, we do not look for any action sequence $(\pi a.a)^*$ but we reduce the search space since we search for sequences of invocations of procedure `manageExecution` with appropriate parameters.

8 Conclusion

Most of existing PMSs are not completely appropriate for very dynamic and pervasive scenarios. Indeed, such scenarios are turbulent and subject to a higher frequency of unexpected contingencies with respect to usual business settings that show a static static and foreseeable behaviour. This paper describes SmartPM, an adaptive PMS that is able to adapt processes thus recovering from exceptions. Adaptation is synthesized automatically without relying either on the intervention of domain experts or on the existence of specific handlers planned in advance to cope with specific exceptions. Space limitation has prevented from including concrete examples of adaptation: interested readers can refer to [11].

Future works aim mostly at integrating SmartPM with state-of-art planners. Indeed, current imple-

mentation relies on the IndiGolog planner, which performs a blind search without using smarter techniques recently proposed to reduce the search space by removing a priori all the possibility surely taking to no solution. The most challenging issue is to convert Action Theories and IndiGolog programs in a way they can be given as input to planners (e.g., converting to PDDL [3]).

Acknowledgement. The author wishes to thank to Giuseppe De Giacomo, Andrea Marrella, Massimo Mecella and Sebastian Sardina, who have contributed to different aspects of the SmartPM development.

References

- [1] K. Andresen & N. Gronau (2005): *An Approach to Increase Adaptability in ERP Systems*. In: *Managing Modern Organizations with Information Technology: Proceedings of the 2005 Information Resources Management Association International Conference*. Idea Group Publishing, pp. 883–885.
- [2] Giuseppe De Giacomo, Raymond Reiter & Mikhail Soutchanski (1998): *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.
- [3] Maria Fox & Derek Long (2006): *Modelling Mixed Discrete-Continuous Domains for Planning*. *Journal of Artificial Intelligence Research* 27, pp. 235–297.
- [4] M. Ghallab, D. Nau & P. Traverso (2004): *Automated Planning: Theory and Practice*. Morgan Kaufmann Publishers.
- [5] Jan Hidders, Marlon Dumas, Wil M. P. van der Aalst, Arthur H. M. ter Hofstede & Jan Verelst (2005): *When are two workflows the same?* In: *CATS '05: Proceedings of the 2005 Australasian symposium on Theory of computing*. Australian Computer Society, Inc., pp. 3–11.
- [6] S. R. Humayoun, T. Catarci, M. de Leoni, A. Marrella, M. Mecella, M. Bortenschlager & R. Steinmann (2009): *The WORKPAD User Interface and Methodology: Developing Smart and Effective Mobile Applications for Emergency Operators*. In: *HCI International 2009: Proceedings of 13th International Conference on Human-Computer Interaction*. Springer. To appear.
- [7] Bartek Kiepuszewski, Arthur H. M. ter Hofstede & Christoph Bussler (2000): *On Structured Workflow Modelling*. In: *CAiSE '00: Proceedings of the 12th International Conference on Advanced Information Systems Engineering*. Springer-Verlag, London, UK, pp. 431–445.
- [8] Robert A. Kowalski (1995): *Using meta-logic to reconcile reactive with rational agents*. *Meta-logics and logic programming*, pp. 227–242.
- [9] Massimiliano de Leoni, Massimo Mecella & Giuseppe De Giacomo (2007): *Highly Dynamic Adaptation in Process Management Systems Through Execution Monitoring*. In: *BPM'07: Proceedings of the 5th International Conference on Business Process Management, Lecture Notes in Computer Science 4714*. Springer, pp. 182–197.
- [10] Massimiliano de Leoni, Massimo Mecella & Ruggero Russo (2007): *A Bayesian Approach for Disconnection Management in Mobile Ad Hoc Networks*. In: *WETICE '07: Proceedings of the 16th IEEE International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises*. IEEE Computer Society, Washington, DC, USA, pp. 62–67.
- [11] Massimiliano de Leoni (2009): *Adaptive Process Management in Highly Dynamic and Pervasive Scenarios*. Computer engineering, SAPIENZA - Università di Roma. Draft Downloadable at <http://www.dis.uniroma1.it/~deleoni/documents/PhD.pdf>.
- [12] R. Milner (1980): *A Calculus of Communicating Systems, Lecture Notes in Computer Science 92*. Springer.
- [13] R. Reiter (2001): *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. MIT Press.
- [14] Sebastian Sardina, Giuseppe De Giacomo, Yves Lespérance & Hector J. Levesque (2004): *On the Semantics of Deliberation in Indigolog—from Theory to Implementation*. *Annals of Mathematics and Artificial Intelligence* 41(2-4), pp. 259–299.