

# Decentralised Process Enactment in a Multi-Perspective Development Environment

Ulf Leonhardt Jeff Kramer Bashar Nuseibeh

Dept. of Computing, Imperial College, 180 Queen's Gate, London SW7 2BZ, UK  
{ul, jk, ban}@doc.ic.ac.uk

Anthony Finkelstein

Dept. of Computer Science, City University, Northampton Square, London EC1V 0HB, UK  
acwf@cs.city.ac.uk

## ABSTRACT

The ViewPoints framework for distributed and concurrent software engineering provides an alternative approach to traditional centralised software development environments. We investigate the use of decentralised process models to drive consistency checking and conflict resolution in this framework. Our process models use pattern matching on local development histories to determine the particular situation (state) of the development process, and employ rules to trigger situation-dependent assistance to the user. We describe how communication between such process models facilitates the decentralised management of explicitly defined consistency constraints in the ViewPoints framework.

## 1 Introduction

Software engineering processes usually involve the participation of a number of people. The more people are involved, the more important becomes the collaboration and communication between the individuals. The different participants will have different views on and assumptions about the problem domain. This necessitates organised interaction including conflict detection and resolution.

In most systems, these centralised control mechanisms are used to check and enforce consistency whenever possible. It has been recognised that such an *eager* approach does not adequately reflect the needs of concurrent and distributed software engineering processes [3].

An alternative is the decentralisation of data storage and consistency control. As a consequence, conflict detection and resolution have to be made based on inter-

action and local, and thus partial, knowledge about the system. The ‘eager’ approach to conflict detection and resolution discussed above is not viable as the complexity introduced by the distribution of control makes it too expensive. Tolerating inconsistencies is often desirable in order to avoid unnecessary restrictions on the development process [15]. Consequently, the focus is shifted from avoidance to management of inconsistencies.

By addressing this issue explicitly (for example [22]), we can devise and support more sophisticated models of cooperation and communication among the members of the development team (see also [8]). A ‘lazy’ approach to consistency detection and enforcement can be taken: synchronise whenever necessary. However, inconsistency management is a complex task. Local agents have to decide what checks to invoke, when to invoke them, and how to keep track of the results. Process support therefore becomes even more crucial.

In this paper we show how fine-grained, decentralised process models can be used to drive conflict detection and resolution. These models are used to guide the developer rather than automate the development process. We describe how the process models initiate and monitor consistency checks in order to gain knowledge about the system under development. These consistency checks are the prime means of coordination for the development of such a system. As a framework for this work we use the ViewPoints approach which has been described in earlier papers [10, 16, 20, 22, 23]. This paper details the process modelling aspects of the framework which are not covered in our earlier publications.

We start our discussion with a brief account of the ViewPoints framework (section 2), followed by a “motivating example” (section 3). Then we describe our process modelling approach (section 4) and its application to decentralised inconsistency management (section 5). In a scenario walk-through at the end of this paper (section 6) we outline how our models of a concurrent software engineering process can guide the human agents involved in such a process.

## 2 ViewPoints

ViewPoints are the building blocks of our framework for supporting distributed software engineering. Each ViewPoint contains an artefact of the development process (for example, a partial specification) together with a thread of development activities concerning this artefact. These are locally managed, and can be characterised as a collection of loosely coupled objects that encapsulate partial knowledge.

A ViewPoint contains knowledge about the notation, tools and strategies it supports—‘method knowledge’; and the results of the application of that knowledge—‘specification knowledge’. A ViewPoint is structured into the following ‘slots’:

*Style* contains the representation scheme in which the partial specification contained in a ViewPoint is represented.

*Work Plan* contains a process model specifying what the ViewPoint user can do and how the user should do it.

*Domain* specifies the part of the modelled system described by the ViewPoint.

*Specification* contains the result of the method user’s activities, i.e. the partial specification.

*Work Record* stores current status, history and rationale of the ViewPoint’s development process. It contains a trace of all the actions in the ViewPoint’s development history.

We use partially instantiated ViewPoints, ViewPoint templates, to specify method knowledge in a reusable way. Thus specific development methods can be implemented as a collection of ViewPoint templates.

### 2.1 Implementation

We have developed a prototype implementation of automated support for the framework called **The Viewer** [20], which has been extended following collaboration with Hewlett-Packard and Siemens [2, 11, 12]. **The Viewer** supports both method design and method use and is therefore both CASE and MetaCASE tool. Template sets supporting methods for requirements engineering and distributed systems design have been implemented [24, 17].

## 3 A Scenario

Our scenario is structured into five steps which, we believe, highlight some of the important issues in concurrent software engineering (see also [7]).

We use data flow diagrams as an example notation. In such a diagram, a node in a graph may be decomposed

in a separate diagram. For such a diagram hierarchy we wish to ensure that decomposition diagrams exist for composite nodes (constraint 1), and that the contextual data flows are the same for nodes and their respective decomposition diagrams (constraint 2). Constraint 1 specifies syntactic completeness, constraints 2 specifies a notion of agreement. These constraints on data flow diagrams form the basis of the consistency checks (for a formal specification see [18]). Here, we will ignore in-ViewPoint constraints and checks.

Composite nodes are shaded grey, primitive nodes are white. We adopt the convention that the domain of each ViewPoint denotes the node of the parent diagram of which it is a decomposition. Only if the domain of the ViewPoint is labelled *top*, a parent node does not exist. In this case, the ViewPoint contains the root node of the data flow diagram hierarchy. The Work Record of the ViewPoint lists the last seven events in the development history of the ViewPoint (the full history is stored).

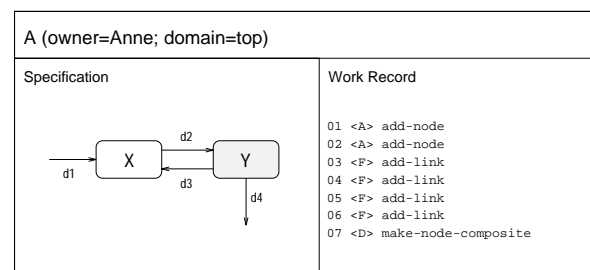


Figure 1: A ViewPoint with a simple data flow diagram in its Specification slot, and a development history listed in its Work Record. ViewPoint A contains the non-primitive node Y for which no corresponding decomposition ViewPoint exists (step 1).

**Step 1** The owner of ViewPoint A, Anne, has developed a top-level data flow model of the system. She has flagged node Y for further decomposition, thus violating global constraint 1 (Figure 1). Subsequently, Anne assigns the responsibility to decompose Y to Bob.

**Step 2** Bob creates a new ViewPoint B the domain of which indicates that it is a decomposition of Y. As the new ViewPoint initially contains an empty specification. Global constraint 2 is not satisfied (Figure 4).

**Step 3** Bob continues developing ViewPoint B by creating a data flow diagram that decomposes Y. When he has finished both ViewPoints satisfy all local and global constraints (Figure 5).

**Step 4** Now, Anne and Bob concurrently develop their ViewPoints as their understanding of the target

system increases. Anne adds an output to Y (d7). Bob does the same but uses a different label (d6). He also adds another output (d9) and renames a third (d4-d8). The result again violates global constraint 2.

**Step 5** While decomposing X Anne realises that its interaction with Y is much more complex than expected. Anticipating the need to restructure ViewPoint A Anne merges the decompositions of X and Y deleting both. Consequently, constraint 1 no longer holds, because the result of the merge is neither a decomposition of X nor Y.

In each of the above steps constraints are temporarily violated. Therefore, there is a need to tolerate such constraint violations in a concurrent development process. In this context, constraints can only be checked and enforced at certain points. Hence consistency must be established by organising the application of the different checks and monitoring their result. As steps 4 and 5 show, this consistency management may be difficult, even in a such a simplified example. Consequently, guidance to users regarding the invocation of consistency checks is necessary. Such guidance must also be tailored to the development method used. In step 4, for example, more frequent checking may be required in order to avoid the accumulation of inconsistencies.

The results of previous consistency checks and all local development activities are stored in individual Work Record slots (see Figure 5, for example). Clearly, this knowledge must be taken into account when deciding when to invoke particular consistency checks.

In the following sections we describe a process modelling approach that addresses these issues. In section 6 we then apply this framework to the scenario presented above.

## 4 Decentralised Process Modelling

In line with the ViewPoints approach, the process modelling framework must support multiple, loosely coupled process models. At run-time there will be no explicit representation of the global process. However, it may be necessary for the method designer to “derive” the local process models from a global model, or to verify certain properties by integrating all local models into a global one.

In this section, we introduce techniques for fine-grained, local process modelling [21] in order to address some of the issues outlined above. We then discuss, how cooperation between process models and other global objectives may be achieved in this context.

### 4.1 Fine-grained local process models

We believe that enactable, fine-grained process models need to address the following issues:

- *Identifying* the current state of the process.
- *Deciding* what course of action is appropriate—taking into account the state of the process.
- *Enacting* the decisions made in the process.

The following sections describe our approach to solving these problems.

#### 4.1.1 Process state

We use *regular expressions* as a concise and easy-to-handle notation to represent finite state machines. Thus, we can make use of a variety of efficient and powerful tools for regular expression handling that are readily available in many programming environments.

Regular grammars define the notion of well-formed input words over some language. For each regular grammar a finite state machine can be constructed that decides whether a given sequence of input characters (word) is well-formed. Thus a regular expression defines an acceptor automaton.

This principle can be applied to ViewPoints by using the sequence of actions and events stored in the *Work Record* as input words for such acceptor automata. Essentially, this amounts to regular expression matching over the development history of a ViewPoint. This process can also be viewed as looking for known patterns of activity in the past of a ViewPoint. The value of an observer predicate associated with a grammar will therefore indicate whether or not the pattern of activity defined by the grammar has been recognised.

#### 4.1.2 Making decisions

The observer predicates and functions<sup>1</sup> define a discrete and finite set of states for the process.

In this setting, we call the mapping of the current state into a course of action a *decision*<sup>2</sup>. Typically, a specific course of action will be appropriate not only for one state but for a set of similar states. We call such a set of states a *situation*. Situations can be defined by logical propositions built from the observer predicates and functions described above.

To express “decision knowledge” in our process modelling framework, we use rules of the general form

**<situation><response>**

The *situation* forms the pre-condition of the rule. That means, the rule fires whenever the current state matches the situation described. *Response* specifies what course of action should be taken as result of the decision made.

<sup>1</sup>certain auxiliary functions, for example `dfd_count_nodes()`

<sup>2</sup>The terms *decision* and *situation* are a variation of the NATURE process meta-model terminology [14], although defined in a different framework.

An extension to these rules would be the addition of post-conditions in the MARVEL-style [4] to support planning activities.

### 4.1.3 Enaction

We distinguish three different types of such responses in decision rules:

- *Informal Guidance.* Here, we assist the user by displaying help texts, video clips, etc. Typically, such assistance would be given in complex and difficult situations.
- *Precise Recommendations.* Specific actions are recommended to the user. In this case, the user is asked to select an action from a limited number of choices. Usually, this applies to well-structured decision problems.
- *Automatic Execution of actions.* This should only occur if the correctness of the decision is reasonably certain and acceptable to the user.

### 4.1.4 Local architecture

Figure 2 shows how local process models for ViewPoints are structured, and how they interact with other components of the ViewPoint. We use the event trace from the work record (shown in abbreviated form as a sequence of tokens) to feed the acceptor automata defined by regular expressions. These, together with the other observer predicates, are matched against the preconditions of the decision rules. If a rule fires, the reaction is enacted on the process.

### 4.1.5 Notation

Our process models consist of *tests* (that is, acceptor automata defined by regular expressions), and *rules* (mapping situations into reactions).

Here, you see the definition of a simple test:

$T_A: \quad . * D [ \neg R ] * \$$	not-successfully- checked-since-D
------------------------------------	--------------------------------------

A test has a short and a long name which enclose a regular expression<sup>3</sup>. In this example, it matches if a **D**-event but no subsequent **R**-event can be found in the local Work Record of the ViewPoint. **D** and **R** are abbreviations for actions or communication events which are defined in the Work Plan.

Rules map a situation into a response to the environment.

<sup>3</sup>We assume familiarity with the basic constructs of regular expressions as used by Lex [5].

Rule: $R_1$
Situation: $T_A \wedge \neg T_B \wedge \neg T_C$
Response: <code>recommend:</code> <code>child-exist-check</code>

The pre-condition, named *situation*, is a logical proposition using tests defined in the process model. Additionally, method specific predicates may also be available in individual templates. The *response* part of the rule describes what should be done when the rule fires. The commands `display`, `recommend`, and `execute` are available to describe such responses.

## 4.2 Communication between process models

We believe that an implementation of the ViewPoints framework can be built on top of a communication system that supports asynchronous message-passing. Therefore we used message-passing between ViewPoints as the basic communication mechanism in the framework<sup>4</sup>.

From the perspective of the communicating process models, a message passing transaction consists of two phases:

1. The source ViewPoint executes an action that sends a message to the destination ViewPoint. The action is appended to the Work Record of the source ViewPoint.
2. Upon receipt, an incoming message is automatically appended to the Work Record of the destination ViewPoint.

We can also use this basic message-passing scheme to build other, more sophisticated communication and cooperation protocols (two-phase locking, for example). In section 5.2 we describe a protocol for two-party consistency checking using message-passing between ViewPoints.

## 5 Managing Consistency

In this section we describe the application of our process modelling framework to consistency management. Our intention to “automate” consistency management necessitates a formal specification of the constraints we want to impose on the system. Such constraints may apply only locally—‘in-ViewPoint’ constraints— or globally—‘inter-ViewPoint’ constraints. Inter-ViewPoint constraints form the basis of coordination between different ViewPoints.

In our framework, consistency checking is decentralised, that is, each ViewPoint checks with the ViewPoint it considers relevant.

<sup>4</sup>We assume some underlying reliable point-to-point communication medium.

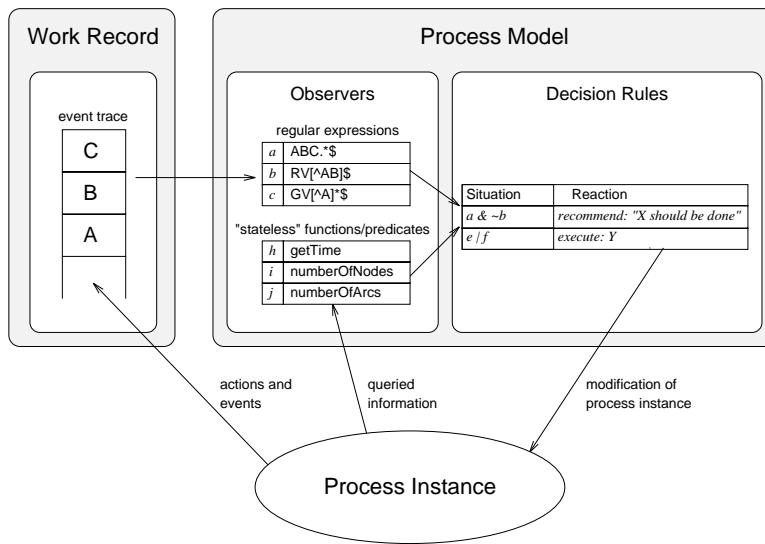


Figure 2: Local architecture

We use local and global consistency constraint to define a desirable state of the decentralised development process. In this sense, we specify a goal for the guidance and assistance that is provided by the process models.

In the remainder of this section, we discuss how local and global consistency constraints can be managed.

### 5.1 Local consistency management

Achieving consistency of the local partial specification contained in some ViewPoint is a necessary subtask of global consistency management. Here, the purpose of the process model is to guide the invocation of local actions and consistency checks.

The local process model sees consistency checks like any other action performed by the ViewPoint. The result of a check is posted to the work record and therefore visible to the process model.

The consistency checks available will vary considerably from template to template. It is the method designer's task to implement specific process models together with the consistency checks required for the different templates.

The consistency checks should be 'fine-grained' because only then the process model can give fine-grained guidance (see also [21]). It is desirable to decompose more complex constraints into independent parts that can be checked separately. Once the checks have been identified, they are integrated into the process model together with the other Work Plan actions.

### 5.2 Global consistency management

In our framework, global consistency management initiates and monitors two-party consistency checks between

ViewPoints. Therefore, we now present a process model driven enactment of two-party consistency checking.

The protocol is based on message passing between ViewPoints as introduced in section 4.2. It assumes that the set of instantiated ViewPoints is constant, that is, a fixed configuration of ViewPoints. We then show how, by self-modification of process models, the general case of varying configurations is addressed.

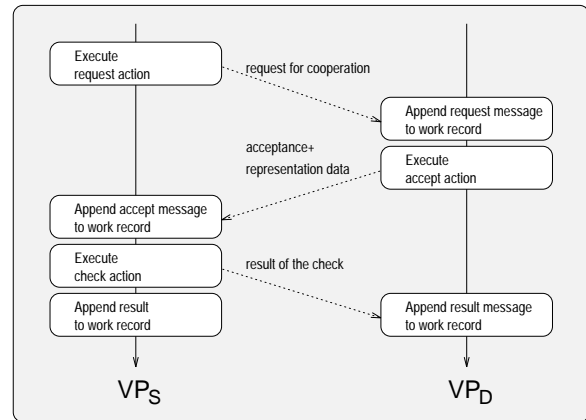


Figure 3: Protocol for two-party checks

**Fixed configurations of ViewPoints** Given that the set of ViewPoints is fixed and all ViewPoints know this, cooperation can be hard-coded into the local finite-state machine process models. A two-party consistency check is carried out by the protocol shown in Figure 3. The actual computation of the check involving the partial specifications contained in both ViewPoints is done

**Variable configurations of ViewPoints** We address this problem by reducing it to the fixed ViewPoints case discussed above. We do so by dynamically modifying the finite state process models as new ViewPoints are created and other ViewPoints are discarded. The results of completeness checks that look for particular ViewPoints are used to update the process model lazily. To cater for a varying set of ViewPoints, parts of the process model must be generic in order to allow for Work Record entries containing ViewPoint identifiers to be processed. The following generalisations of the process modelling architecture are necessary:

- Regular expression templates<sup>5</sup> with a ViewPoint identifier as a parameter have to be used to handle communication messages. Instances of such templates behave like the ‘ordinary’ regular expression discussed above. For example:

$T_g(v)$ : <code>.*q(v)[^P]*\$</code> child-located
---

Here  $q(v)$  is the generic event necessitating the abstraction of the regular grammar.

- To handle dynamically created regular expressions, rule templates are introduced. We restrict ourselves to one ViewPoint identifier as argument. For example:

Rule:	$R_4(v)$
Situation:	$T_g(v)$
Response:	recommend: child-agrees-check-request-(v)

Such a rule would be instantiated and deleted together with the relevant regular grammars concerning a specific ViewPoint. There is an instance of this rule for each known ViewPoint with which communication can take place.

- We also want to be able to express statements like: *If all checks have succeeded do X*. Technically, this requires the expression of a universal quantification at some point in our model. Here quantification over instances of the same regular expression template plays this role. For example:

Rule:	$R_7(v)$
Situation:	$(\forall v).T_i(v)$
Response:	display: “All children known have been checked successfully”

<sup>5</sup>These must not be confused with ViewPoint templates.

The pre-condition of this rule is satisfied if all instances of the template  $T_i(v)$  match.

- Sometimes it is not necessary to know the identity of the other party involved when responding to a communication event. In this case, we ignore the address part of a communication message. For example:

$T_f$ : <code>.*U[~r(*)]*\$</code> commissioned-vp-not-sighted-yet
--

Here  $r(*)$  matches any instance of  $r(v)$ . Such regular expressions have one instance only, and do not require special treatment as far as the rules are concerned.

On the basis of such generic tests and rules we believe that communication of arbitrary and evolving configurations of ViewPoints can be handled.

### 5.3 Coordinating the checks

The framework described for two-party consistency checking necessitates cooperation among ViewPoints. Therefore the method designer developing process models has to look at the system in its entirety rather than at a specific ViewPoint template. This global perspective plays an important role by guiding and verifying the design of the local process models. The task of composing and decomposing process models, however, is non-trivial and requires tool support.

It is difficult to give general rules governing how consistency checks should be coordinated. Again, a recommended *sequence* of checks could be described. The notion of state shown has to be modified because remote actions influencing the state of a constraint will, in general, not be observable. Therefore, we have to resort to more heuristic measurements of the state. For example the *age* of a check (that is, the number of local actions and events since the last successful check) could be interpreted as reflecting the probability that the constraint still holds.

## 6 The Scenario Revisited

We now demonstrate our process modelling approach by applying it to the scenario described in section 3.

### 6.1 DFD ViewPoints

The ViewPoints used in this scenario each contain a simple data flow diagram (DFD). Consequently, all ViewPoints are instances of the same ViewPoint template which defines a DFD technique. A simple example showing the Specification and Work Record slots of one such ViewPoint is shown in figure 1.

**Checks** To detect violations of the constraints listed, we use local and two-party consistency checks. They are derived from the consistency constraints outlined in section 3.

**Two-party checks** We now decompose the global constraints into two-party constraints.

- **parent-exists-check** detects whether the ViewPoint has a parent node. The check also succeeds if the domain is labelled *top*.
- **child-exist-check** takes all the composite nodes in the local diagram and tries to find the corresponding decomposition ViewPoints. As a side-effect of this check, the process model is modified by adding and/or removing generic actions, tests (that is, acceptor automata defined by regular expressions) and rules that control the interaction with decomposition ViewPoints.
- **parent-agrees-check-do-(v)** is a generic action performing a check for agreement between the local DFD and the DFD provided by *v*. Agreement in this case means the contextual data flows of the child node must match those of the parent. These must also be parametrised.
- **child-agrees-check-do-(v)** performs the same operation as above, but in the opposite direction.

**Process Model** Part of the work plan is the process model consisting of regular grammars  $T_x$  and rules  $R_i$ . Actions, rules and tests can have one parameter, then they are *generic*. Generic parts of the work plan are instantiated and deleted as side-effects of existence checks (for example, **parent-exists-check**).

The regular expressions  $T_x$  as presented here operate on sequences of action tokens. The action tokens for static events are fixed, and the tokens for generic events are dynamically allocated. Here are all the tokens used in the regular expression examples in the scenario:

A	add-node
D	make-node-composite
F	add-link
G	remove-link
H	rename-link
O	parent-exist-check
Q	parent-exist-check-succeeded
R	child-exist-check-succeeded
T	child-exist-check-failed
U	commission-viewpoint
a(v)	parent-agrees-check-request-(v)
i(v)	parent-agrees-check-requested-(v)
m(v)	parent-agrees-check-succeeded-(v)
q(v)	child-located-me-(v)
r(v)	parent-located-me-(v)

Except for the parametrised tokens necessitated by these dynamic allocations, the syntax and semantics of the regular expressions follow the usage by the lexical analyser generator **Lex** [5].

The rules  $R_i$  uses the regular expressions  $T_x$  as predicates in their pre-conditions. There are three alternatives that can be used in the action part of such rules:

**do:<action>** automatically executes the given *action*.

**recommend:<action>** suggests that *action* be enacted by the user.

**display:<text>** gives informal guidance by displaying the help message *text*.

## 6.2 Process enactment at work

We now outline how our process modelling framework applies to the scenario described in section 3. We have selected steps 2 and 3 to discuss how guidance should be provided to the user. The reader may also find it useful to consult section 3 on page 2 to recall some of the details.

**Step 2** The ViewPoints A and B do not agree with each other (Figure 4). Potentially, both A and B can discover the inconsistency. However, given that the Specification of B is empty, B is unlikely to initiate the check and may refuse to cooperate with A on this matter. ViewPoint A waits for some acknowledgement message from the decomposition ViewPoint because A initiated B's creation. Effectively, the checking of global check 2 is therefore suspended as long as B remains empty. The process model of B, however, will advise Bob to elaborate ViewPoint B, thus making it non-empty. The checking of global constraint 1 at ViewPoint A would detect the presence of B and update the local process model.

**Example** The invocation of checks by Anne or Bob can be inhibited by the following tests and rules:

$T_D$ :	$\sim.\{0,10\}\$$	underdeveloped
$T_f$ :	$.*U[\sim r(*)]*\$$	commissioned-vp-not-sighted-yet

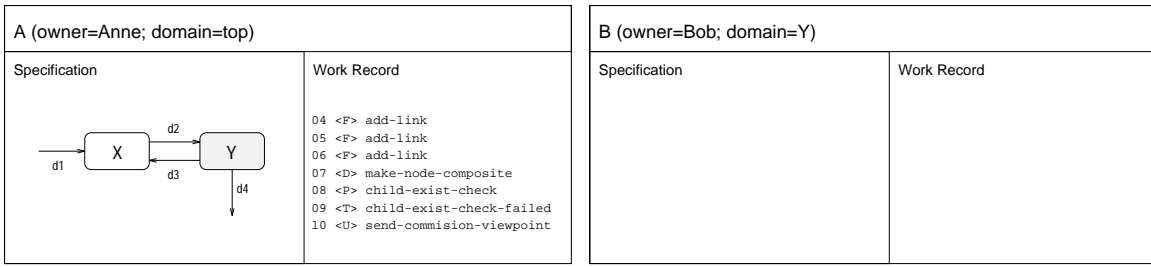


Figure 4: Now a decomposition ViewPoint for Y exists, but the inputs and outputs do not match (step 2).

<b>Rule:</b> $R_2$ <b>Situation:</b> $\neg T_D$ <b>Response:</b> display: "Viewpoint should be developed further"
<b>Rule:</b> $R_3$ <b>Situation:</b> $T_f$ <b>Response:</b> display: "do not enact child-exist-check or child-agrees-check-request because ViewPoint creation commissioned but not acknowledged"

In this situation  $R_2$  advises Bob to perform some additional local development before considering any non-local checks.  $R_3$  tells Anne to wait until the ViewPoint to be created by Bob reports its existence to ViewPoint A.

**Step 3** Now, all constraints are satisfied (Figure 5). All global checks are eventually initiated by A or B with their result communicated to the work record of both ViewPoints. The checks carried out separately, however, do not guarantee that both ViewPoints are consistent, even at any single point in time. To achieve this would require two-party checks to be performed as atomic transactions.

**Example** The initiation of the agreement check between A and B is recommended by these tests and rules:

$T_g(v)$ : <code>.*q(v)[^P]*\$</code> child-located
$T_E$ : <code>.*Q[^0]*\$</code> parent-probably-exists
<b>Rule:</b> $R_5$ <b>Situation:</b> $T_E$ <b>Response:</b> recommend: parent-agrees-check-request-(v)
<b>Rule:</b> $R_4(v)$ <b>Situation:</b> $T_g(v)$ <b>Response:</b> recommend: child-agrees-check-request-(v)

For ViewPoint A,  $T_g(B)$  succeeds when ViewPoint B has been identified. Consequently, Anne is advised to check the agreement with the child whose existence is now reasonably certain.  $T_E$  tests for a recent success of the check for the existence of B's parent. If this is the case, Bob can also initiate the agreement check with ViewPoint B's parent, A.

## 7 Related Work

Iida et al. describe the use of communicating tasks defined by regular expressions to model cooperative software development processes [13]. Unlike our work however, their system focuses on programmed processes and strict enaction, rather than guidance and tolerated deviation.

Ben-Shaul and Kaiser proposed an approach towards "modelling and enaction of inter-group collaboration among independent, autonomous, and, possibly, pre-existing processes" [6]. They presented an "international alliance" metaphor to define collaboration in terms of 'treaties' which are enacted at 'summits'. Conceptually, our consistency constraints fulfil the role of such treaties, and the application of a consistency check can also be interpreted as a summit between two ViewPoints. Our approaches differ, in that we focus on intra-group collaboration whereas Ben-Shaul and Kaiser address cooperation among groups.

Jarke et al. describe decision-oriented, logically centralised process models [14]. They propose process models dedicated to guidance which use pattern matching to identify situations that can be mapped into guidance, a clear similarity to our work. They also use process models to structure process traces, a task fulfilled (in a less flexible way) by the Work Records in our framework. Jarke et al. address the issue of deriving guidance from these process traces. We have identified this problem in our framework and are currently investigating possible solutions.

Balzer [3] describes a mechanism for managing inconsistencies employing "pollution markers" to identify constraints that have been violated, which can then be either avoided or resolved. We have generalised and de-



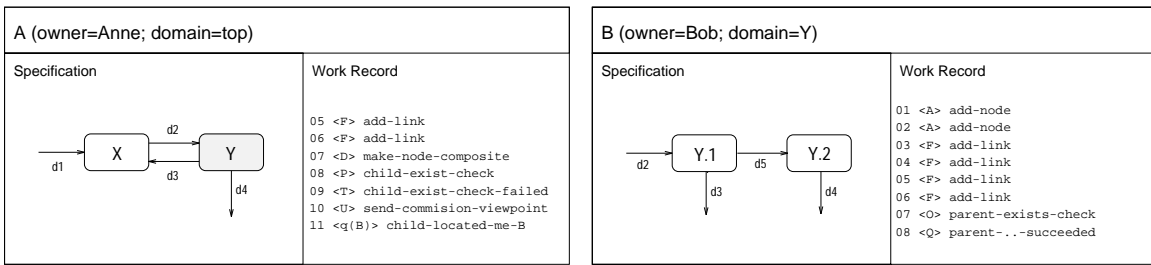


Figure 5: The specification is now consistent with respect to the formulated constraints (step 3).

centralised these concepts in our approach. In his work, Balzer also uses enforced consistency constraints, which we consider problematic to adopt in our framework.

Narayanaswamy and Goldman advocate “lazy” consistency as basis for cooperative software development [19]. They focus mainly on preventing conflicts of updates of shared artefacts, where we concentrate on conflict detection and resolution. In their setting, dependencies between products are defined by a single explicit relation. In contrast, our framework allows for many different kind of such dependencies to be defined in terms of consistency constraints. This gives us more detailed information for the resolution of inconsistencies.

## 8 Further Work

We have identified three major areas in need of further work. These are: guidance; support for evolution; process specification support. We see the role of process modelling as providing guidance rather than automation. In this case the way in which the guidance is framed and the means by which the guidance is actually delivered are critical. In this paper we have concentrated on process observation and decision making, however the practical utility of the techniques described will depend on advances in guidance. Some indication of our approach to this problem is given in [9]. We have indicated above some of the problems of evolution in an environment in which ViewPoints are created and (occasionally) destroyed, and we have outlined our solution to these problems. We feel that this area is in need of further work and to this end we will be exploring other grammar-based schemes. We have, as yet, little experience of how to actually arrive at process models of the form we have presented. Compositional process model design requires attention, in particular derivation of local process models from higher level cooperation policies. There are a number of improvements we would like to make to our implementation, of which the most pressing is integration with **The Viewer**; also some improvements to the user interface we provide to the process modelling capability are required.

## 9 Conclusions

This paper has examined the application of process modelling techniques to the problem of consistency management. The approach proposed in this paper is based on constraint checks derived from a static notion of consistency.

We have developed an architecture for communicating local process models which is fully decentralised. In this respect, we believe, it reflects the nature of development processes in multi-perspective environments (exemplified by the ViewPoints framework). Global and local consistency checking is driven by local process models employing regular grammars and rules. This we achieve by decomposing and distributing global strategies and protocols for cooperation and communication. We have presented a communication mechanism based on message passing between ViewPoints, and consequently their local process models. We have also shown that protocols for two-party consistency checking can be built on top of this communication layer.

Our process models are ‘fine-grained’, and therefore capture the level of detail which we believe to be essential for adequate guidance. The granularity of process models in our framework critically depends on the granularity of consistency checks. Therefore we also favour fine-grained consistency constraints.

The application of the proposed process modelling architecture to the scenario has demonstrated the process-model driven consistency management as the innovative feature presented in this paper. We have also developed a prototype implementation of our process modelling framework which we have used to validate the ideas described.

## 10 Acknowledgements

We would like to gratefully acknowledge the constructive comments of Michael Goedicke. This work was partly funded by the UK Department of Trade and Industry (DTI) as part of the ESF project, by the UK EPSRC VOILA project, and by the European Union (HCM, ISI, ESPRIT BRA PROMOTER).

## REFERENCES

- [1] In A. van Lamsweerde and A. Fugetta, editors, *Proceedings of the 3rd European Software Engineering Conference (ESEC '91)*, volume 550 of *LNCS*, Milan, Italy, October 1991. Springer-Verlag.
- [2] L. Ballesteros. Using ViewPoints to Support the FUSION Object-Oriented Method. M.Sc. Thesis, Department of Computing, Imperial College, London, UK, September 1992.
- [3] R. Balzer. Tolerating inconsistency. In *Proceedings of the 13th International Conference on Software Engineering*, Austin, Texas, May 1991. IEEE CS press.
- [4] N. Barghouti and G. Kaiser. Scaling up rule-based software development environments. In van Lamsweerde and Fugetta [1], pages 380–395.
- [5] Bell Telephone Laboratories, Inc., Murray Hill, New Jersey. *UNIX programmer's manual*, seventh edition, 1983. Volume 2.
- [6] I. Ben-Shaul and G. Kaiser. A paradigm for decentralized process modeling and its realization in the OZ environment. In *Proceedings of the 16th International Conference on Software Engineering*, pages 179–188, Sorrento, Italy, May 1994. IEEE CS press.
- [7] S. Easterbrook, A. Finkelstein, J. Kramer, and B. Nuseibeh. Coordinating Distributed ViewPoints: the anatomy of a consistency check. *International Journal on Concurrent Engineering: Research and Applications, Special issue on conflict management*, 2(3), 1994.
- [8] A. Finkelstein, D. Gabbay, A. Hunter, J. Kramer, and B. Nuseibeh. Inconsistency handling in multiperspective specifications. In *IEEE Transactions on Software Engineering*, August 1994.
- [9] A. Finkelstein and J. Kramer. TARA: Tool assisted requirements analysis. In *Conceptual Modelling, Databases & CASE: an integrated view of information system development*. McGraw Hill, 1991.
- [10] A. Finkelstein, J. Kramer, B. Nuseibeh, M. Goedicke, and L. Finkelstein. ViewPoints: A Framework for integrating multiple Perspectives in System Development. *International Journal of Software Engineering and Knowledge Engineering*, 2(1):31–58, March 1992.
- [11] P. Graubmann. The HyperView Tool Standard Methods. REX technical report REX-WP3-SIE-008-V1.0, Siemens, Germany, July 1990.
- [12] P. Graubmann. The Petri Net Method ViewPoints in the HyperView Tool. REX technical report REX-WP3-SIE-021-V1.0, Siemens, Germany, January 1992.
- [13] H. Iida, K. Mimura, K. Inoue, and K. Torii. Hakoniwa: Monitor and navigation system for cooperative development based on activity sequence model. In *2nd International Conference on the Software Process*, pages 64–74, Berlin, Germany, 1993.
- [14] M. Jarke, K. Pohl, C. Rolland, and J. Schmitt. Experience-Based Method Evaluation and Improvement: A Process Modeling Approach. NATURE Report Series 94-15, ESPRIT Project 6353, RWTH Aachen, Germany, 1994.
- [15] J. Kramer. CASE Support for the Software Process: A Research Viewpoint. In van Lamsweerde and Fugetta [1].
- [16] J. Kramer and A. Finkelstein. A configurable framework for method and tool integration. In *European Symposium on Software Development Environments and CASE*, volume 509 of *LNCS*, pages 233–257, Königswinter, Germany, June 1991. Springer-Verlag.
- [17] Fui Kien Lai. CORE in The Viewer. M.Sc. Thesis, Department of Computing, Imperial College, London, UK, September 1993.
- [18] U. Leonhardt, A. Finkelstein, J. Kramer, and B. Nuseibeh. Decentralised process enactment. Technical Report 95/5, Department of Computing, Imperial College, London, UK, January 1995.
- [19] K. Narayanaswamy and N. Goldman. “Lazy” Consistency: A Basis for Cooperative Software Development. In *Proceedings of CSCW'92*, pages 257–264, Toronto, Canada, 1992. ACM press.
- [20] B. Nuseibeh and A. Finkelstein. Viewpoints: A vehicle for method and tool integration. In *Proceedings of the International Workshop on Computer-Aided Software Engineering (CASE '92)*, pages 50–60, Montreal, Canada, July 1992. IEEE CS press.
- [21] B. Nuseibeh, A. Finkelstein, and J. Kramer. Fine-grain process modelling. In *Proceedings of the 7th International Workshop on Software Specification and Design (IWSSD-7)*, pages 42–46, Redondo Beach, California, December 1993. IEEE CS Press.
- [22] B. Nuseibeh, J. Kramer, and A. Finkelstein. Expressing the relationships between multiple views in requirements specification. In *Proceedings of the 15th International Conference on Software Engineering*, pages 187–196, Baltimore, Maryland, May 1993. IEEE CS press.
- [23] B. Nuseibeh, J. Kramer, and A. Finkelstein. A framework for expressing the relationships between multiple views in requirements specification. In *IEEE Transactions on Software Engineering*, volume 20, pages 760–773. IEEE CS Press, October 1994.
- [24] T. Thanitsukkarn. The Constructive Viewer. M.Sc. Thesis, Department of Computing, Imperial College, London, UK, September 1993.

Papers related to the ViewPoints framework can be found on <ftp://dse.doc.ic.ac.uk/dse-papers/>.