

Workflow Modeling using Procllets

W.M.P. van der Aalst^{1,2}, P. Barthelmess², C.A. Ellis², and J. Wainer^{2,3}

¹ Department of Technology Management, Eindhoven University of Technology, P.O. Box 513, NL-5600 MB, Eindhoven, The Netherlands. w.m.p.v.d.aalst@tm.tue.nl

² Department of Computer Science, University of Colorado at Boulder, Campus Box 430, Boulder, CO 80309-0430, USA. {barthelm, skip}@colorado.edu

³ Department of Computer Science, State University of Campinas, Caixa Postal 6176, 13083-970, Campinas - SP, Brazil wainer@dcc.unicamp.br

Abstract. The focus of traditional workflow management systems is on control flow *within one* process definition. The process definition describes how a single case (i.e., workflow instance) in isolation is handled. For many applications this paradigm is inadequate. Interaction between cases to support communication and collaboration is at least as important. This paper introduces and advocates the use of interacting *procllets*, i.e., light-weight workflow processes. By promoting interactions to first-class citizens it is possible to model complex workflows in a more natural manner. In addition, the expressive power and flexibility are improved compared to the more traditional workflow modeling languages.

1 Introduction

In the last decade many workflow management systems have become available [20]. These systems allow for the explicit representation and support of business processes and avoid the need to re-code applications every time a business process changes. As the workflow paradigm continues to infiltrate organizations that need to cope with complex administrative processes, it is becoming apparent that the available workflow management systems have difficulties dealing with the increasingly dynamic and inter-organizational nature of today's business processes [29]. As we will argue in this paper, one of the core problems of the current generation of workflow languages and tools is the focus on isolated case-based processes.

Perspectives that are relevant for workflow modeling and workflow execution are : (1) *control-flow* (or process) perspective, (2) *resource* (or organization) perspective, (3) *data* (or information) perspective, (4) *task* (or function) perspective, and (5) *operation* (or application) perspective. (These perspectives are similar to the perspectives given in [17].) In this paper we primarily focus on the control-flow perspective. This does not imply that the other perspectives are less relevant. However, the problems addressed in this paper are mainly related to the control-flow perspective. In traditional workflow management systems, the control-flow perspective of a workflow is described by one *workflow process definition* (also called workflow schema). A workflow process definition specifies which tasks need to be executed and in what order (i.e., the routing or control flow). A *task* is an atomic piece of work. Workflow process definitions are instantiated for specific *cases* (i.e., workflow instances). Examples of cases are a request

for a mortgage loan, an insurance claim, a tax declaration, an order, or a request for information. Since a case is an instantiation of a process definition, it corresponds to the execution of concrete work according to the specified routing.

Today's workflow management systems predominantly focus on the control-flow *within one* process definition. This assumes that a workflow process can be modeled by specifying the *life-cycle of a single case in isolation*. For many real-life applications this assumption is too restrictive. As a result, the workflow process is changed to accommodate the workflow management system, the control-flow of several cases is artificially squeezed into one process definition, or the coordination amongst cases is hidden inside custom built applications. Consider for example an engineering process of a product consisting of multiple components. Some of the tasks in this engineering process are executed for the whole product, e.g., the task to specify product requirements. Other tasks are executed at the level of components, e.g., determine the power consumption of a component. Since a product can have a variable number of components and the components are engineered concurrently, it is typically not possible to squeeze this workflow into one process definition. This is a direct consequence of the fact that, in most workflow management systems, the degree of parallelism is fixed in workflow process definition, i.e., it is not possible to concurrently instantiate selected parts of the workflow process a variable number of times. Using iteration one can instantiate parts a variable number of times. However, this results in the sequential execution of inherently parallel tasks.

To solve these problems, we propose an approach based on *proclets*, *performatives* and *channels*. Proclets are light-weight processes. Typically, a proclet represents only one aspect or one element of the whole workflow. The interaction between proclets is modeled explicitly, i.e., proclets can exchange structured messages, called performatives, through channels. By adopting this approach the problems related to purely case-based processes can be avoided.

The remainder of this paper is organized as follows. First, we motivate our approach by clearly identifying the problems encountered when modeling the reviewing process of a conference. Then we present the framework which is based on Petri nets [25, 26] and inspired by concepts originating from object-orientation [8, 27], agent-orientation [18], and the language/action perspective [14, 34–36]. In Section 4, we model the reviewing process using our framework. Finally, we compare the framework with existing approaches and conclude with our plans for future research.

2 Motivating Example: Organizing a Conference

The process of selecting papers for a conference presents features that challenge existing modeling languages. In brief, the goal of this process is to select some papers out of a normally larger set, based on different criteria (e.g., quality, minimum and maximum number of papers). After a set of people is invited and accepts to act as program committee members, a call for papers is issued to prospective authors. These authors submit papers that are then subject to review by peers (invited by pc members) and finally a selection is made. A very brief and abstract sequence of steps would be:

- Invite program committee (PC) members: these are going to be responsible for the management of reviews.
- Issue a call for papers: this step announces the upcoming conference and asks for submissions.
- Receive the submissions and check them: papers are accepted up until a deadline. Submissions are checked for consistency with conference standards, and so on.
- Distribution: after the submission deadline, each paper is assigned to multiple PC members. These PC members will be responsible for finding reviewers for the papers assigned to them. The goal is to obtain at least a minimum number of reviews, by different people, for each paper.
- Review: reviewing starts with the assignment of a reviewer by a PC member. The paper is made available for the reviewer and after a while a review is produced.
- Selection: after the reviews are completed, the papers are compared and ranked according to the reviewers recommendations and other subjective criteria (e.g., desired number of papers, acceptable quality threshold).
- Notification: authors are notified either of acceptance or rejection of their papers. In case of acceptance, final versions have to be sent in by the authors.
- Publication: finally, the final versions are assembled and sent for publication.

The process is complicated by a series of factors, that we list in a non-exhaustive way:

- Prospective PC members and reviewers may accept or reject the invitation to join the committee and to review one or more papers, respectively. Replacements for those that rejected need to be found.
- Reviewers can fail to return the reviews on time. They may either declare that they are not going to meet the deadline or simply forget about the deadline altogether. As a result, some of the papers may lack enough reviews to allow their fair evaluation.
- The distribution of the papers takes into account varied criteria, such as the number of submitted papers, the number of available PC members, preferences and areas of expertise of PC member's groups, balance of the work load assigned to each PC member (as compared to their availability) and so on. The decision of how to split the papers needs to take into account the whole set of available papers and can not be performed in isolation.
- Selection is a yet more subjective task. Once again, this task can only be performed on the whole set of available papers. Paper quality needs to be gauged against the quality of all the remaining papers, to a certain extent, or at least to a set of related papers. If two or more papers discuss the same topic from different or opposing perspectives, this needs to be taken into account. Other factors, such as the minimum and maximum number of desired papers also influence this task.

A modeler faces many problems translating these requirements. A first basic question is what is to be considered the case⁴ - the submission, the review, the "empty slot" in the conference that one wants to fill with a good quality paper, or is the case the whole set of slots? The choice of each of these as the unit of modeling causes problems. A closer examination of the tasks of such a process reveals that while some of the

⁴ Workflow instance.

tasks operate on each submitted paper individually, others are based on the whole set of papers, and others still are related to individual reviews. While tasks such as *receiving* and *checking*, for example, can be conducted at the level of individual papers, tasks such as *distribution to program committee members* and *final selection* are based on the whole set of papers. *Review* tasks operate on each of the multiple reviews that are produced for each paper.

The class diagram (Figure 1) shows that different tasks rely on information that is at different levels of aggregation - some of the tasks operate at the conference level, that groups all papers, others at the paper level, and others yet at the lower level of review. The choice of any of the possible aggregations as the main one introduces problems whenever we have to deal with the others. One of the major obstacles is, therefore, how to conciliate these multiple perspectives into one model.

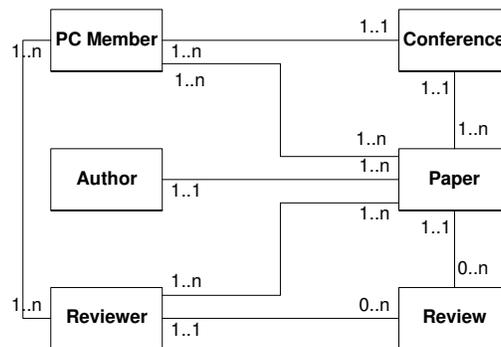


Fig. 1. Review process class diagram.

Lacking the power to express the difference in aggregation, most workflow management systems force one to depict the process at an arbitrarily chosen level (usually the paper level), essentially ignoring the issues that are relevant at the conference and to some extent at the review levels. The resulting models present some important shortcomings:

- The models are artificially flattened, being unable to account for the mix of different perspectives that coexist in the real process. Given that workflow enactment is guided by what is specified by the process model, the missing perspectives will have to be handled and coordinated manually by the users themselves, without further help from the system.
- Batch-oriented tasks are typically not supported. Batch-oriented tasks are those that are based on groupings of lower aggregation elements, e.g., the whole set of papers during distribution and selection, or the set of reviews for a paper, while deciding if enough reviews are available. In other words, it is usually not possible to handle higher aggregation tasks using lower aggregation instances, e.g., conference level tasks within a paper level case.

- Handling lower aggregation tasks within higher aggregation ones is also hard in most languages. Launching and then synchronizing a variable number of reviews (lower aggregation) from a paper centered case, for instance, can not be usually represented in most languages.
- The interactions with the environment are usually abstracted away as well. An important aspect in many processes is the exchange of messages between the entities. Reviewers, for instance: receive invitations to review papers; respond to them by either accepting or rejecting; must be notified of approaching deadlines; send their completed reviews or sometimes send notification of inability to complete reviews. These interactions need to be reflected in the process model, but usually are not.

Conference review is not an atypical example, in the sense that one encounters similar problems very frequently in other areas as well. We next list just a few of the innumerable real world examples where interactions between instances and different levels of aggregation play a strong role:

- In engineering processes: processing of subparts may impact one or more higher level components that make use of this common subpart. Conversely, decisions at the higher level component processes may have an impact on subpart processes. For example, an approaching deadline for a higher level component may cause interruption of the process of certain subparts.
- In software development: software modules are composed of submodules, which in turn may be composed of sub-submodules and so on. Considerations at higher or lower levels of aggregation may influence other levels, e.g., the discovery of some specification flaw at a lower level may have a ripple effect on a variable number of modules at all other levels. Code that is shared by multiple versions also introduces interactions that are hard to model.
- Processing of insurance claims: some claims may refer to the same accident. Even though they may start out as independent instances, at some point in time it is desirable that all related claims be merged so that a uniform decision can be reached.
- Hiring new people: some job applications are received in response to an advertised open position. Candidates have to be evaluated and ranked with respect to each other. Again, the interactions between the applications are most relevant. Some applications are sent in independently of open positions. In this situation, interesting applications may cause a position to be specially created. Again, there is a strong interaction between two perspectives, that of the application and of the position. Sometimes one is the central one, sometimes it is the other.

In summary, we see as limitations of current modeling formalisms 1) the fact that one is usually forced to choose to represent a process at one single level, even when a problem space consists of entities with varying aggregations, 2) that the interactions with the environment can not be made explicit, even though a subjacent model may be (and usually is) assumed.

3 Framework

The examples given in the previous section show that today's workflow management systems typically have problems dealing with workflow processes that are not entirely

case-oriented. Squeezing the control flow of a workflow process into a single process definition often results in unreadable workflow specifications where essential parts of the control flow are hidden inside custom made application software. In fact, there are plenty of examples where the workflow process is changed in order to fit the workflow management system. Clearly, this is undesirable: Workflow technology should support rather than dictate work processes.

Inspired by these problems, we have developed a new framework for modeling workflows. This framework is based on *proclets*. A procllet can be seen as a lightweight workflow process equipped with a knowledge base containing information on previous interactions. One can think of procllets as objects equipped with an explicit lifecycle (in the object-oriented sense) [8, 27] or active documents (i.e., documents aware of tasks and processes) [19]. Procllets interact via *channels*. A channel is the medium to transport messages from one procllet to another. The channel can be used to send a message to a specific procllet or a group of procllets (i.e., multicast). Based on the properties of the channel different kinds of interaction are supported, e.g., push/pull, synchronous/asynchronous, and verbal/non-verbal. In order for procllets to find each other, there is a *naming service*. The naming service keeps track of registered procllets and can be queried by any procllet. The concepts *procllet*, *channel* and *naming service* constitute a framework for modeling workflow processes (see Figure 2).

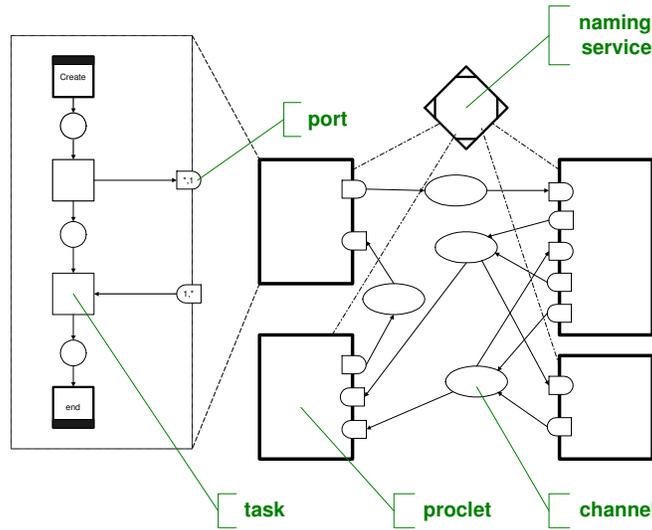


Fig. 2. Graphical representation of the framework.

Compared to existing workflow modeling languages, complex case-based workflow definitions describing the control flow of an entire process are broken up into smaller interacting procllets, i.e., there is a shift from control to communication. The framework is based on a solid process modeling technique (Petri nets [25, 26]) extended with

concepts originating from object-orientation [8, 27], agent-orientation [18], and the language/action perspective [14, 34–36].

In the remainder of this section we present the four main components of our framework: *proclefs*, *channels*, *naming service*, and *actors*.

3.1 Proclefs

A *proclef class* describes the life-cycle of *proclef instances*. A proclef class can be compared to an ordinary workflow process definition or workflow type [17]. The class describes the order in which tasks can or need to be executed for individual instances of the class, i.e., it is the specification of a generic process. Proclef instances can be created and destroyed, and are executed according to a class specification. At any moment a proclef instance has a state. When no confusion is possible we will simply use the term “proclef” instead of “proclef class” and/or “proclef instance”.

To define proclefs, we introduce some preliminaries including some basic Petri net concepts and terminology.

To specify proclef classes, we use a graphical language based on *Petri nets*. Petri nets are an established tool for modeling and analyzing workflow processes [1, 2, 5, 11, 12]. On the one hand, Petri nets can be used as a design language for the specification of complex workflows. On the other hand, Petri net theory provides for powerful analysis techniques which can be used to verify the correctness of workflow procedures [25, 26]. A (classical) Petri net is a directed bipartite graph with two node types called *places* and *transitions*. The nodes are connected via directed *arcs*. Connections between two nodes of the same type are not allowed. Places are represented by circles and transitions by rectangles. A place p is called an *input place* of a transition t iff there exists a directed arc from p to t . Place p is called an *output place* of transition t iff there exists a directed arc from t to p . At any time a place contains zero or more *tokens*, drawn as black dots. The *state*, often referred to as marking, is the distribution of tokens over places. The number of tokens may change during the execution of the net. Transitions are the active components in a Petri net: they change the state of the net according to the following *firing rule*:

- (1) A transition t is said to be *enabled* iff each input place p of t contains at least one token.
- (2) An enabled transition may *fire*. If transition t fires, then t *consumes* one token from each input place p of t and *produces* one token in each output place p of t .

Petri nets can move from one state to another by firing enabled transitions. A state s is *reachable* if there is a sequence of transition firings which leads from the current state to state s . A Petri net in a given state is *safe* if for any reachable state no place contains multiple tokens, i.e., the number of tokens per place is limited to 1. A Petri net in a given state is *live* if for any reachable state s and for any transition t it is possible to reach a state from s such that t is enabled. A transition t is called *dead* if there is no reachable state enabling t . Reachable, safe, live, and dead are standard concepts which can be found in any textbook on Petri nets [25].

In this paper, we use a specific subclass of Petri nets. This subclass corresponds to the so-called class of sound WF-nets [1, 2].⁵ A WF-net has source and sink transitions: A source transition has no input places and a sink transition has no output places. Every node (i.e., place or transition) is on a path from some source transition to some sink transition. Moreover, any WF-net is connected, i.e., the network structure cannot be partitioned in two unconnected parts. A WF-net becomes activated if one of the source transitions fires. In the remainder we assume that a WF-net becomes activated only once (single activation). A WF-net is called *sound* if and only if the following requirements are satisfied:

- (1) *safeness*: Each state reachable under the single activation assumption is safe.
- (2) *proper completion*: Firing one of the sink transitions empties the net, i.e., after firing a sink transition no tokens are left.
- (3) *completion option*: From any reachable state it is possible to reach a state which enables one of the sink transitions, i.e., termination is always possible.
- (4) *dead transitions*: there are no dead transitions.

These four requirements are quite reasonable in the context of workflow management: It should always be possible to terminate properly, there should be no parts which cannot be activated, and, since the WF-net will model one proclat instance, it should be safe. Soundness can be verified using state-of-the-art analysis techniques [1, 2]. Based on these techniques we have developed a workflow verifier called Woflan [33].

Most workflow modeling languages primarily focus on control flow inside one process definition and (partly) abstract from the interactions between process definitions, i.e., coordination is limited to the scope of the process definition and communication and collaboration are treated as second-class citizens. Therefore, our framework explicitly models interactions between proclats. The explicit representation of interactions is inspired by the *language/action perspective* [36, 35] which was introduced in the field of information systems by Flores and Ludlow [14] in the early 1980's and is rooted in *speech act* theory [28]. In contrast to traditional views of "data flow" the language/action perspective emphasizes what people do while communicating; how they create a common reality by means of language and how communication brings about a coordination of their activities. The need for treating interaction as first-class citizens is also recognized in the agent community [18]. Emerging agent communication languages such as KQML [13] demonstrate this need.

Inspired by these different perspectives on interaction, we use *performatives* to specify communication and collaboration among proclats. A performative is a message exchanged between one sender proclat and one or more receiver proclats. A performative has the following attributes:

- (1) *time*: the moment the performative was created/received.
- (2) *channel*: the medium used to exchange the performative.
- (3) *sender*: the identifier of the proclat creating the performative.
- (4) *set of receivers*: the identifiers of the proclats receiving the performative, i.e., a list of recipients.

⁵ For the readers familiar with WF-nets: For notational convenience we omit the unique source and sink place used in [1, 2].

- (5) *action*: the type of the performative.
- (6) *content*: the actual information that is being exchanged.

The role of these attributes will be explained later. At this point, it is important to note the action attribute. This attribute can be used to specify the illocutionary point of the performative. The five illocutionary points identified by Searle [28] (assertive, directive, commissive, declarative, expressive) can be used to specify the intent of the performative. Examples of typed performatives identified by Winograd and Flores are request, offer, acknowledge, promise, decline, counter-offer or commit-to-commit [36]. In this paper, we do not restrict our model to any single classification of performatives (i.e., a fixed set of types). However, at the same time we stress the importance of using the experience and results reported by researchers working on the language/action perspective.

Proclefs combine performatives and sound WF-nets. A *proclef class PC* is defined as follows:

- (1) *PC* has a *unique name*. This name serves as a unique identification of the class which we will refer to as *class_id*.
- (2) *PC* has a *process definition* defined in terms of a *sound WF-net*. The transitions correspond to *tasks* and the places correspond to *state conditions*.
- (3) *PC* has *ports*. Ports are used to interact with other proclefs. Every port is connected to one transition.
- (4) Transitions can send and receive *performatives* via ports. Each port has two attributes: (a) its *cardinality* and (b) its *multiplicity*. The cardinality specifies the number of recipients of performatives exchanged via the port. The multiplicity specifies the number of performatives exchanged via the port during the lifetime of any instance of the class.
- (5) *PC* has a *knowledge base* for storing these performatives: Every performative sent or received is stored in the knowledge base.
- (6) Tasks can query the knowledge base. A task may have a *precondition* based on the knowledge base. A task is enabled if (a) the corresponding transition in the WF-net is enabled, (b) the precondition evaluates to true, and (c) each input port contains a performative.
- (7) Tasks connected to ports have *post conditions*. The post condition specifies the outcome of the task in terms of performatives generated for its output ports. The performatives which are generated may depend upon information obtained from the *naming service* (i.e., proclef identifiers).

A proclef class is a generic definition, i.e., it does not describe the behavior and properties of a specific proclef. Proclef (instances) are created by instantiating the proclef class and have a *unique identification* which we will refer to as *proc_id*. Note that for a concrete proclef all elements, i.e., the process definition, knowledge base, ports, and tasks, are instantiated. For example, the tokens in the WF-net specifying the process definition refer to one proclef instance, i.e., tokens of different proclef instances are *not* merged into one WF-net. (Recall that a sound WF-net is safe.) Moreover, each proclef instance has its own private knowledge base. However, proclef instances can *share* performatives

with all other instances of the same class. This means that part of the knowledge base is *public* and part of the knowledge base is *private*. The public part is identical for all instances of the class, i.e., effectively this part resides at the class level. The private part exclusively resides at the instance level. Whenever a performative is sent or received, the corresponding procelet decides whether it should be stored in the public or in the private part.

A performative has by definition one sender, but can have multiple recipients. The sender is always represented by a *proc_id*, i.e., the identifier of a procelet instance. However, the list of recipients can be a mixture of *proc_id*'s and *class_id*'s, i.e., one can send performatives to both procelet instances and procelet classes. A performative sent to a procelet class is received by all procelet instances of that class.

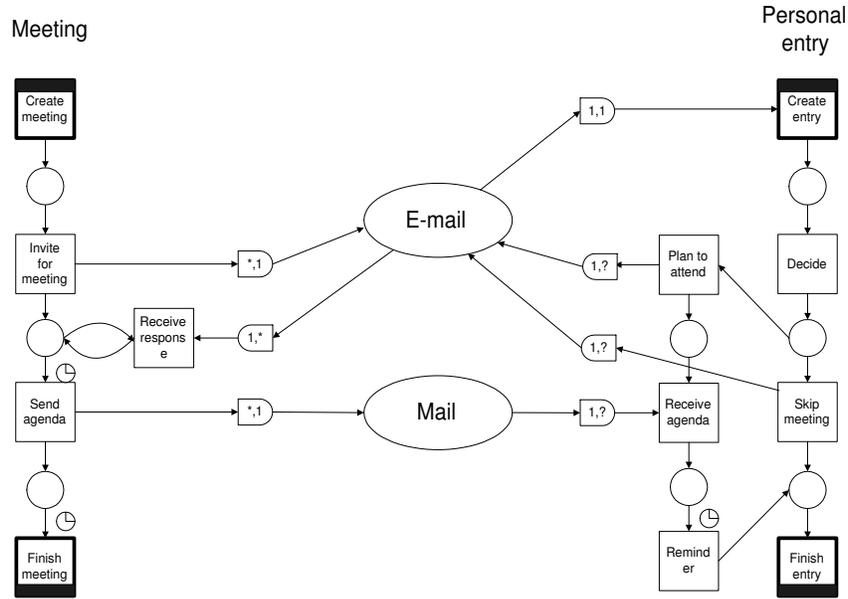


Fig. 3. Example of two procelet classes: *Meeting* and *Personal entry*.

To illustrate the framework we use the example shown in Figure 3. There are two procelet classes. Both classes are used to organize meetings. Procelet class *Meeting* is instantiated once per meeting. Procelet class *Personal entry* is instantiated for every potential participant of a specific meeting. The instance of class *Meeting* first sends an invitation to all potential participants. The *proc_id*'s are used to multicast the invitation performative to a specified set of instances of class *Personal entry*. Note that the cardinality of the port connected to task *Invite for meeting* is denoted by a star *. This star indicates that the invitation is sent to an arbitrary number of potential participants, i.e., the performative has multiple recipients. We will use * to denote an arbitrary number of recipients, + to denote at least one recipient, 1 to denote precisely one recipient, and ?

to denote no or just a single recipient. Performatives with no recipients are considered not to have occurred, i.e., only performatives with a positive number of recipients are registered in the knowledge base. The multiplicity of the output port connected to task *Invite for meeting* is denoted by the number 1. This means that during the lifetime of an instance of class *Meeting* exactly one performative is sent via this port. The invitation performative is sent through the channel *E-mail* (The role of channels is explained in Section 3.2). The performative creates a procelet for each recipient, i.e., creation task *Create entry* is triggered. Creation tasks are depicted by squares with a black top. The input port connected to *Create entry* has cardinality 1 and multiplicity 1. Every input port has by definition cardinality 1, i.e., from the perspective of the receiving procelet there is only one procelet receiving the performative. Input ports connected to a creation task (i.e., a source transition) have by definition a multiplicity of 1 or ?: An instance can be created only once. Since there is just one creation task in *Personal entry*, the multiplicity is 1. After an instance of the class *Personal entry* is created a decision is made (task *Decide*). Based on this decision either task *Skip meeting* or *Plan to attend* is executed. In both cases a performative is sent to the instance of the procelet class *Meeting*. The performative is either a confirmation (*Plan to attend*) or a notification of absence (*Skip meeting*). Note that each instance of the class *Personal entry* sends such a performative. These performatives are sent through channel *E-mail*. Note that the ports connected to *Plan to attend* and *Skip meeting* both have cardinality 1 (i.e., one recipient) and multiplicity ? (one performative is sent via one of the two ports). Task *Receive response* is executed once for every “confirmation/notification of absence” performative. Therefore, the corresponding port has multiplicity *. After some time, as indicated by the clock symbol [2], task *Send agenda* is executed. In this small example we assume that all potential participants respond before this time-out occurs. *Send agenda* generates one performative: the agenda of the meeting. This performative is sent to all procleets that confirmed the invitation. This performative has multiple recipients, i.e., the cardinality of the corresponding output port is *. Since the agenda is sent only once the multiplicity is 1. The procleets that confirmed the invitation receive the agenda (task *Receive agenda*) and a timer for the task *Reminder* is set. Finally, all procleets are destroyed by executing the finishing tasks *Finish meeting* and *Finish entry*. The finishing tasks (i.e., sink transitions) are depicted by squares with a black bottom.

3.2 Communication Channels

Communication channels are used to link procleets. Channels transmit messages containing performatives from sending procleets to receiving procleets. There are many different categories of channels defined by channel properties such as medium type, reliability, security, synchronicity, closure, and formality. These properties are briefly explained:

– *Medium Type*

This can be point-to-point or broadcast, or some form of limited multicast. Recall that performatives can be sent to an individual procelet instance (point-to-point), a set of procleets (multicast), or an entire procelet class (broadcast). Common media include postal mail, telephone, and electronic mail. Different media satisfy different communication requirements. We are also concerned with media of face-to-face

communication such as sound waves of spoken voice, gestures, and body language. The framework presented in this paper, assumes that there is only one sending procler. However, there are situations where a group effort results in a single performative (e.g., orchestral performances). In fact there are many examples that could not be accomplished by a single person or procler (e.g., collaborations modeled as single acts such as lifting a heavy object). Such group efforts can be modeled by introducing a so-called proxy procler. This procler coordinates and consolidates the group effort.

– *Reliability*

Some channels are very reliable; some are unreliable. For some electronic channels, we assume that the technology is robust, and that error detection and retransmission are implemented at lower layers of the communication protocols. In this case, we need not be concerned with these details in our higher level modeling. Thus, channels built upon TCP/IP are more reliable than those built upon UDP. A problem of dial-in data channels in some lesser developed countries is that the channel (the phone lines) are inherently unreliable. Thus, sometimes the data gets sent, and sometimes not. Similar unreliability is sometimes exhibited by postal services. A different channel available from the postal service is registered mail, where the cost of mailing a letter is higher, and the reliability is also higher.

– *Security*

At times the content of a performative is considered to be quite valuable and secret. In such cases, the transmission should be via highly secure channels. In electronic transmission, encoding and encryption are sometimes used to implement secure channels.

– *Synchronicity*

This is concerned with the time delay of message delivery and acknowledgment. Some channels are used for real time communications in which each party expects to get rather immediate feedback from recipient parties. This requires synchronous channels. Face-to-face spoken conversation falls into this category. In other cases, the expectation is that the recipient will not instantaneously receive the message content. In the case of an asynchronous channel, the sender usually is not waiting for an immediate response. For example, when email is sent, there is usually no expectation of immediate response. When a UNIX talk session is initiated, there is expectation of immediate response.

– *Closure*

Channels can be classified as open or closed channels. When a channel is open, the sender does not know exactly who, and how many recipients are connected. When a channel is closed, the exact identity of all recipients is specified in advance. A radio broadcast, and a notice posted on a bulletin board are respectively examples of synchronous and asynchronous communications in which the medium is open because the senders do not exactly know who are their recipients.

– *Formality*

Some channels convey much more formality in the messages delivered than others. Performatives can be very formally specified, or can be informal and flexible. Generally, business letters are much more formal than chat rooms. A scheduled meeting with a rigid agenda is much more formal than a casual conversation over coffee. A

careful record is kept of formal channel transmissions, whereas informal channels are usually not recorded; they are "off the record."

Clearly, channel properties and performative types are closely related, i.e., for a given performative certain properties are appropriate others are not. For example, for the performative "You are fired!" a point-to-point, reliable, secure, synchronous, closed, and formal channel is most appropriate.

3.3 Naming service

All interaction is based on proclat identifiers (*proc_id*'s) and class identifiers (*class_id*'s). These identifiers provide the handles to route performatives. By sending a performative using a *class_id*, all instances of the corresponding class receive the performative. Only if a proclat knows the *proc_id*'s of the recipients of the performative, it is able to communicate with specific proclats. In many situations the sending proclat does not know the *proc_id*'s of all receiving proclats. Therefore, we introduce the concept of the *naming service*. The naming service keeps track of all proclats and can be queried to obtain *proc_id*'s. There are many ways to implement such a naming service. Consider for example the services provided by the object request brokers developed in the context of CORBA. In this paper, we only consider the desired functionality and abstract from implementation details (e.g., distribution of the naming service over multiple domains).

The naming service provides the following primitives: *register*, *parent*, *child*, *update*, *unregister*, *query*, and *forward*.

The function *register* is called by the proclat the moment it is created. Therefore, the execution of one of the create tasks (i.e., source transitions) coincides with the execution of the *register* primitive. The primitive has the following parameters: *creator* (i.e., the *proc_id* of the calling proclat), *time* (i.e., the time the function is called), *class name* (i.e., the name of class of the created instance), *owner* (i.e., the identity of the actor responsible for the proclat) and *attributes* (i.e., the characteristic properties of the created proclat) and returns a new unique *proc_id*. The *proc_id* is returned by the naming service in order for the proclat to know its own identity. Proclats can be created by other proclats. Consider for example Figure 3. The create task *Create entry* is triggered by a performative sent by a *Meeting* proclat. The performative is created by the task *Invite for meeting*. This implies that the task *Invite for meeting* already registered the new *Personal entry* proclat. The new proclat is already registered by the meeting proclat because the meeting proclat needs a handle to the newly created proclat. Since proclats can be created by other proclats, there are parent-child relationships. The functions *parent* and *child* can be used to navigate through the naming service. Both functions have a *proc_id* parameter. The *parent* function returns a *proc_id* (if any) and the *child* function returns a set of *proc_id*'s.

The proclat attributes registered in the naming service describe the essential characteristics, e.g., role and group attributes, links to actors, etc. The set of attributes is not fixed and may vary from one class to another. During the life cycle of the proclat these attributes may change. The function *update* with parameters *proc_id* and *attributes* can be used to change existing or add new attributes.

Based on the attributes, proplets can query the naming service using the function *query*. The function has one parameter describing a Boolean expression in terms of attributes and returns a set of *proc_id*'s, i.e., all proplets satisfying the expression.

Entries in the naming service can be removed using the function *unregister*. Executing a finish task (i.e., a sink transition in the WF-net) results in a call to *unregister*. Function *unregister* has one parameter: The *proc_id* of the proplet to be destroyed.

Sometimes there is a need to merge proplets. Consider for example two proplets corresponding to the same traffic accident. If two police officers file a report on the same traffic accident, two proplets are created. If after executing some steps it turns out that both proplets correspond to the same traffic accident, then it does not make sense to execute the remaining tasks for both proplets. Therefore, we propose to merge the two proplets by destroying one of them and redirecting all performatives to the remaining one. For this purpose we propose the function *forward*. This function has two *proc_id* parameters: one for the destroyed proplet and one for the remaining proplet. As a result of calling this function, all performatives intended for the destroyed proplet are redirected to the remaining proplet.

3.4 Actors

Proplets have *owners*. Owners are the actors responsible for the proplet. Actors can be automated components, persons, organizations (e.g., shipping department), or even whole companies. Owners are specified at proplet registration time and this information is kept by the naming service (see Section 3.3). Ownership can be transferred by updating the naming service information.

The owner will sometimes be the executor of proplet tasks him or herself - in the example of Figure 3, for instance, the owner of the personal entry will most probably be the one that will perform the tasks, essentially the decision of attending or skipping the meeting. Roles may be specified for each task, in which case the executor can be different from the owner. We assume that the usual role resolution mechanisms [21] are employed in this latter case.

We propose to model as *external proplets* those actors (in the broad sense of the word) that interact with proplets in a more complex way. *External proplets* are useful to model those interactions that go beyond the simple model assumed by the usual role mechanism, e.g., when a request for service may be either accepted, rejected or counter-proposed. *External proplets*, as the name implies, represent entities that are outside of the scope of process proper, whereas *internal proplets* are those under the control of the workflow system enactment service. Both types of proplets are modeled in a similar way - by describing expected interactions with other proplets. Detailed examples of both *internal* and *external proplets* are present in Section 4.

4 Example Revisited

We now revisit the conference review process, this time using proplets. The multiple perspectives of conference, paper and review that were identified in Section 2 as being

one of the obstacles for representation are taken into account and integrated into a seamless model. The resulting model has a much broader scope than the ones usually found in the literature. In particular, interactions with the environment are made explicit.

The model is composed of six proclets, with well defined interfaces, that correspond to the class diagram entities previously presented (Figure 1). Three of the proclets correspond to *internal proclets* (Figures 4, 5, 6) and the other three are *external* (Figures 7, 8, 9).

The *Conference* procllet groups tasks that act upon or require access to the set of all submitted papers, e.g., the distribution among PC members and final selection of papers. For each *Conference* procllet, there will exist many related *Paper* procllets - one instance per paper. Each *Paper* procllet will in turn be associated to some *Review* procllets. There will be as many *Review* procllets as there are reviews. The multiple instances of *Paper* and *Review* procllets directly reflect the multiple cardinality of the relationships between conference, paper and review as shown in the class diagram (Figure 1). *Author*, *PC member* and *Reviewer* are external procllets and specify the details of the interactions between these actors and the internal procllets.

We now analyze in more detail the *Conference* procllet (Figure 4). The first few tasks in this procllet deal with the staffing of the program committee (PC). *Invite PC member* sends out a multicast message to prospective PC members. These invitations will either be accepted or rejected. In case of rejection, a new round of invitations can take place. Notice that here the responses to the illocutionary act *invite* are explicitly included in the model. The single multicasted invitation will be responded to asynchronously by the persons that were invited, so the tasks *Accepted* and *Rejected* are enabled in a loop and receive multiple messages, one at a time, as indicated by the cardinality 1, and multiplicity * of the associated ports.

This part of the model illustrates the need and use of knowledge bases. The task *Replace rejected* should obviously only fire if one or more *rejection* performatives were received. *Replace rejection* has a pre-condition that queries the knowledge base and only allows firing if at least one *rejection* has been received. Similarly, as soon as a certain number of PC members have accepted the invitation, the pre-condition for the task *Call for papers* will enable it to fire.

After the committee is staffed, a call for papers is issued, multicasted to many prospective authors. In practice the recipients of this multicasted performative will be mailing lists and individuals whose identities are stored in some database. Once again, the responses will be received one by one, in separate asynchronously generated messages. *Receive paper* therefore is enabled in a loop that receives the submissions and that sends a performative that creates new instances of *Paper* procllets. The result is that there will eventually exist as many *Paper* procllet instances as there are submissions.

The *Distribution* task of the *Conference* procllet corresponds to the decision making as to which paper should be handled by which PC members. Once this decision is made, a performative informs each *Paper* procllet (Figure 5) about the identity of assigned PC members. The *Paper* procllet, in turn, generates a second multicast, this time to create as many *Review* procllets (Figure 6) as needed (one for each assigned PC member). The performative to each *Review* procllet informs about the identity of the responsible PC member. This illustrates one basic design principle - work is distributed through

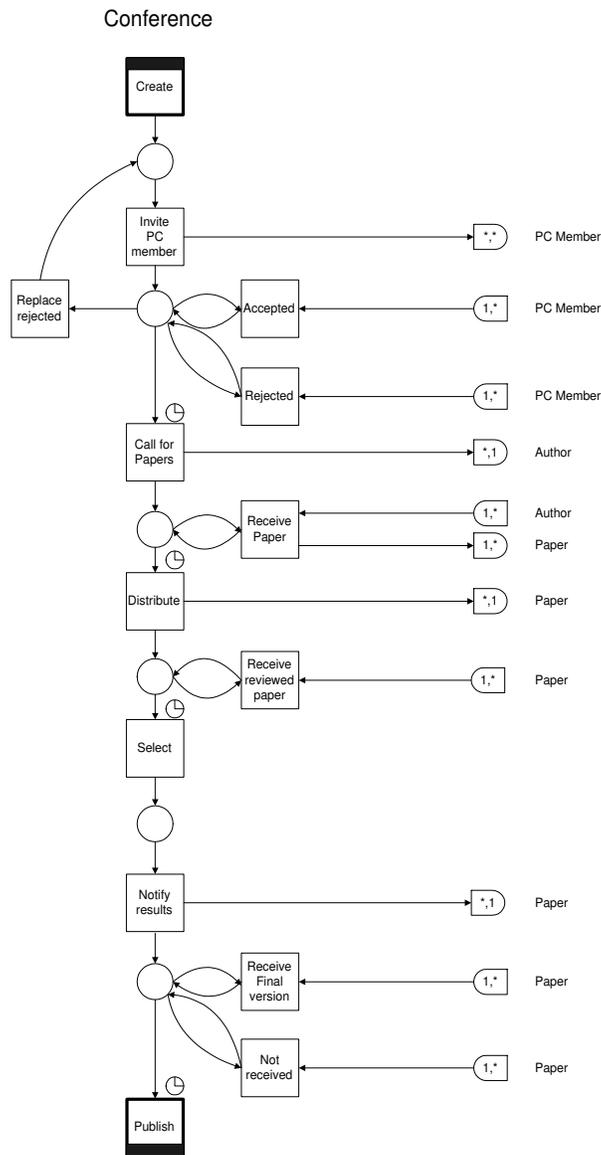


Fig. 4. Conference procelet.

the procleets in such a way that each procelet deals only with tasks that are at the same aggregation level. The *Conference* procelet, for instance, groups tasks that operate on the whole set of submitted papers, while *Paper* procleets handle work at the individual paper level. *Review* procleets group tasks that pertain to each of the multiple individual reviews a paper has.

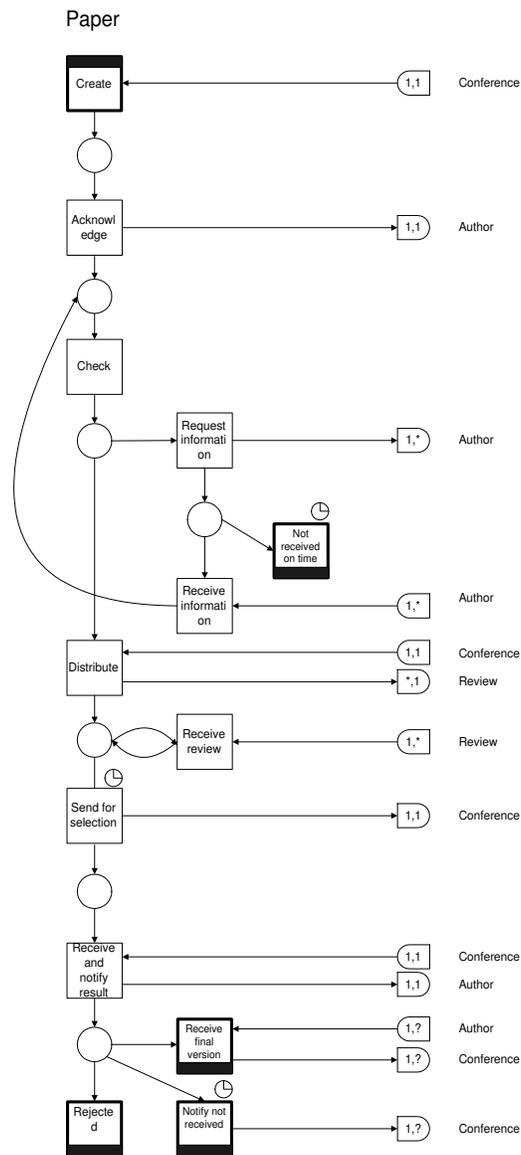


Fig. 5. Paper procleets.

Back in the *Conference* procelet, *Selection* decides whether papers should be accepted or rejected based on their relative merits. The final decision is multicasted to the Paper procleets, that notify the authors and then wait for the reception of final versions of those papers that were accepted.

spect to the conference at some point in time. *External procllets* are lightweight in the sense that they do not imply that the environment conduct business in that specific manner, only that it is compatible with the specified communication behavior. Typically, *external procllets* do not correspond to an executing object, and usually just reflect the fact that in the environment there is an actor that can be expected to behave according to some communication protocol. *External procllets* allow us to make these assumptions explicit, making them visible and verifiable, through inspection and/or simulation.

Initially, authors receive the call for papers (or hear about it from a friend), submit papers (or not), receive acknowledgments and provide requested information (if any) until the submission deadline. These possible interactions are modeled by the *Author* procllet (Figure 7). From the point of view of this procllet, there are no explicit constraints on the order in which these messages will be generated/received. Note that one author can submit multiple papers for the same conference. Therefore, acknowledgments, submissions, etc., can be interleaved.

In a similar way, *PC member* and *Reviewer* procllets (Figures 8 and 9) model expected interactions. An important difference is that explicit responses from these actors are expected, specifically regarding the invitations to join the process. While authors will not typically inform the PC that they are not interested in submitting papers, acceptance or rejection of invitations on the part of prospective PC members and reviewers have a direct impact on the process - acceptance implies commitment to perform required work, and rejection causes actions to find replacements.

Another aspect worth examining is the way by which reviewers are invited to review a paper. Each of the multiple instances of the *Review* procllet will execute task *Request reviewer*, asking the responsible PC member to assign a reviewer. After an assignment is received, the *Review* procllet requests service from this prospective reviewer, by sending a *request* performative. In case of rejection, the procllet itself manages the request for a replacement. These steps are repeated until either a willing reviewer is found or time runs out.

Notice that the model presented here includes aspects that cannot be represented by other modeling languages. In particular:

- The different perspectives, corresponding to the three different level of aggregation, conference, paper, and review, that were identified in the class diagram (Figure 1) are explicitly represented.
- The transition between these different levels of aggregation are cleanly specified as communication between procllets.
- Launching of variable number of instances of lower aggregation elements, and their synchronization - grouping and ungrouping - can be easily and clearly represented.

As motivated in the introduction, traditional workflow management systems are unable to deal with these issues. As a result, the workflow process is changed to accommodate the workflow management system, the control-flow of several cases is artificially squeezed into one process definition, or the coordination amongst cases is hidden inside custom built applications. These unsatisfactory “patches” can be avoided by adopting the framework presented in this paper. The framework also encourages broadening the scope of what is represented, making explicit some of the usually hidden assumptions:

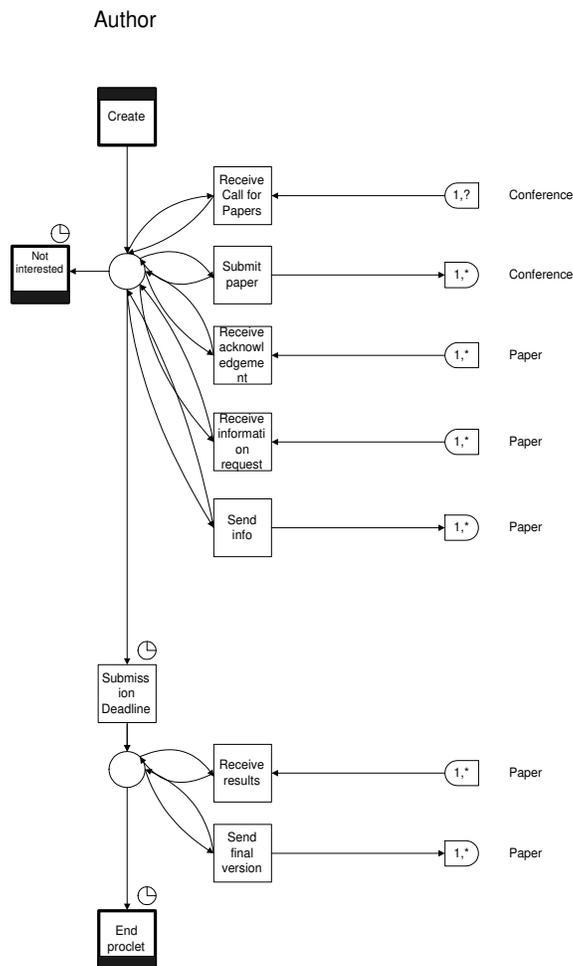


Fig. 7. Author procelet.

- *External procleets* can be used to represent actors that are part of the environment. These are typically omitted from models, which makes them harder to verify.
- *Performatives* offer a mechanism to more precisely model message content. Speech-act theory can be used to clarify and regulate the semantics of interactions.
- *Channels* offer a way to explicitly represent (and eventually support at enactment) different media and their attributes.

It is also important to realize that, even though much more is represented, the resulting model is composed of small modules with clearcut interfaces to the environment. Furthermore, these modules have a one to one correspondence with the entities of the class diagram, which can, therefore, be used as a guideline for procelet development. This is usually not the case in existing modeling languages: models often are mono-

PC Member

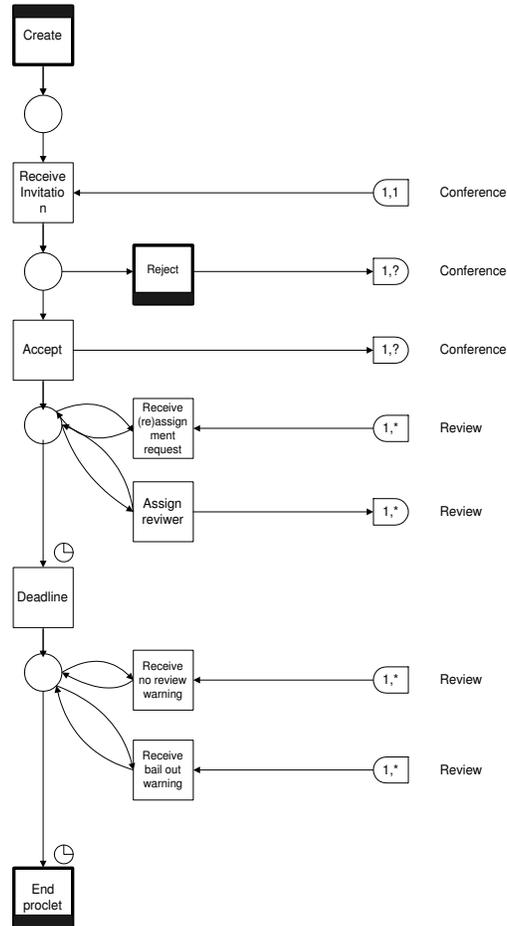


Fig. 8. PC member proclat.

lithic and large; there is usually no close connection between the resulting models and the problem space, as mapped, for instance, by a class diagram. All these features come in addition to the full expressive power of Petri nets, a formalism that has proven to be specially adequate for representing processes in general and workflows in particular.

5 Related work

Petri nets have been proposed for modeling workflow process definitions long before the term “workflow management” was coined and workflow management systems became readily available. Consider for example the work on Information Control Nets, a

Reviewer

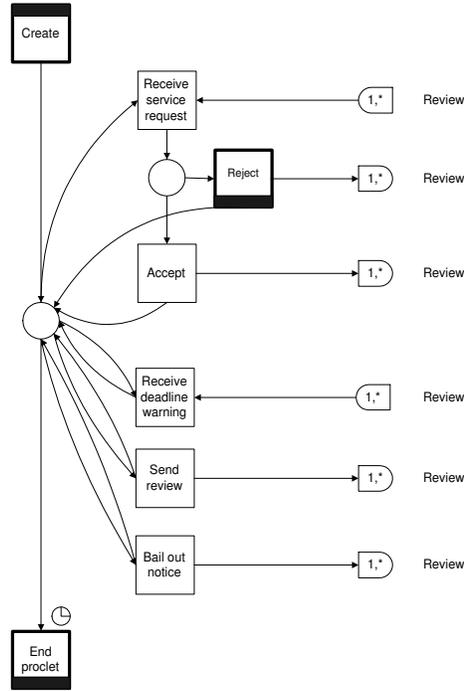


Fig. 9. Reviewer procler.

variant of the classical Petri nets, in the late seventies [11, 12]. Since then many workflow models and languages have been developed ranging from approaches based on other formal models such as state charts [22] to the vendor-specific diagramming techniques used in the many commercial workflow management systems available today. Workflow models described in the literature focus on various aspects (cf. [29]) such as transactional concepts [15], flexibility [24], analysis [1, 2], and cross-organizational workflows [3, 4], etc. Any attempt to give a complete overview of these models is destined to fail. Therefore, we only acknowledge the work that extended workflow models to accommodate the problems identified in Section 2.

Zisman [37] presents a paper refereeing example that involves Petri-nets and allows multiple instantiation of the reviewer net.

In [7], batch-oriented tasks were proposed, i.e., a task is executed for multiple instances at the same time. To allow for the batch-oriented tasks, independent cases need to be synchronized. As an example, consider the task of selecting papers for a conference (task *select* in the example): All papers are considered at the same time. The need to deal with the batch-oriented clustering of instances was also recognized in [10]. A similar extension proposed in [9] were so-called *multi-tasks*. A multi-task is a task in

a process which can be instantiated an arbitrary number of times. The multi-task completes the moment that the corresponding task is completed for each instance, or for some number of those instances (the quorum). A similar mechanism has been implemented for the Regatta system by Fujitsu [30]. In this system, multiple instances are created, according to the number of actors available to perform them. In Spade-1, a process-centered software engineering environment (PCSEE), it is possible to instantiate dynamically the same activity a variable number of times, generating different execution threads for the activity, called *active copies* [6].

The idea to promote interactions to first-class citizens was proposed in different settings. For example, in the context of the language/action perspective [14, 34–36], Action Technologies developed a workflow tool [31] where each step in the process is characterized by four phases: preparation, negotiation, performance and acceptance. The transition from one phase to another is mainly driven by interactions between actors. In the more systems-oriented domains there have also been some proposals for inter-process communication. Consider for example Opera [16], a process support system kernel (i.e., a rudimentary workflow management system), which supports the interaction between different processes.

The language-action perspective is also employed in the context of agent technology [34]. Speech-acts form the basis for performatives in agent interaction languages, e.g. KQML [13]. The use of agents for implementing workflow systems is explored, e.g., in the Bond multi-agent system [32]. Petri-nets are used in Bond as an intermediate representation of workflows [23].

Some of the ideas presented in this section have been adopted by our framework: batch-oriented operation, multi-tasks, and inter-process communication can be handled easily by the framework. In addition, the framework employs concepts such as performatives, channels, ports, knowledge bases, naming services, and the rigor of a Petri-net basis which allows for various forms of analysis and a straightforward and efficient implementation.

6 Conclusion

In this paper, we presented a framework which advocates the use of interacting proplets, i.e., light-weight workflow processes communicating by exchanging performatives through channels. As was demonstrated in this paper, the framework can solve many of the traditional modeling problems resulting from the case-oriented paradigm.

In the future, we plan to explore the relation between channels and performatives. We are also compiling a list of interaction patterns. In our view, the interaction between proplets typically follows a number of well-defined patterns, e.g., a request performative is followed by an accept or reject performative. Finally, we plan to build a prototype to support the framework. This prototype will be used to support the reviewing process of the ACM biannual Siggroup conference following the model described in this paper.

References

1. W.M.P. van der Aalst. Verification of Workflow Nets. In P. Azéma and G. Balbo, editors,

- Application and Theory of Petri Nets 1997*, volume 1248 of *Lecture Notes in Computer Science*, pages 407–426. Springer-Verlag, Berlin, 1997.
2. W.M.P. van der Aalst. The Application of Petri Nets to Workflow Management. *The Journal of Circuits, Systems and Computers*, 8(1):21–66, 1998.
 3. W.M.P. van der Aalst. Interorganizational Workflows: An Approach based on Message Sequence Charts and Petri Nets. *Systems Analysis - Modelling - Simulation*, 34(3):335–367, 1999.
 4. W.M.P. van der Aalst. Process-oriented Architectures for Electronic Commerce and Interorganizational Workflow. *Information Systems*, 24(8):639–671, 2000.
 5. N.R. Adam, V. Atluri, and W. Huang. Modeling and Analysis of Workflows using Petri Nets. *Journal of Intelligent Information Systems*, 10(2):131–158, 1998.
 6. S. Bandinelli, M. Braga, A. Fuggetta, and L. Lavazza. Cooperation support in the spade environment: a case study. In *Proceedings of the Workshop on Computer Supported Cooperative Work, Petri nets, and Related Formalisms (14th International Conference on Application and Theory of Petri Nets)*, Chicago, June 1993. <ftp://ftp-se.elet.polimi.it/dist/Papers/ProcessModeling/CSCWPN93.ps>.
 7. P. Barthelmeß and J. Wainer. Workflow systems: a few definitions and a few suggestions. In N. Comstock and C.A. Ellis, editors, *Proceedings of the Conference on Organizational Computing Systems - COOCS'95*, pages 138–147, Milpitas, California, September 1995. ACM Press.
 8. G. Booch, J. Rumbaugh, and I. Jacobson. *The Unified Modeling Language User Guide*. Addison Wesley, Reading, MA, USA, 1998.
 9. F. Casati, S. Ceri, B. Pernici, and G. Pozzi. Conceptual modeling of workflows. In *Proceedings of the OOER International Conference*, Gold Coast, Australia, 1995.
 10. F. Casati, P. Grefen, B. Pernici, G. Pozzi, and G. Snchez. Wide workflow model and architecture. Technical report, Dipartimento di Elettronica e Informazione, Politecnico di Milano, 1996. <http://dis.sema.es/projects/WIDE/Documents/ase30.4.ps.gz>.
 11. C.A. Ellis. Information Control Nets: A Mathematical Model of Office Information Flow. In *Proceedings of the Conference on Simulation, Measurement and Modeling of Computer Systems*, pages 225–240, Boulder, Colorado, 1979. ACM Press.
 12. C.A. Ellis and G.J. Nutt. Modelling and Enactment of Workflow Systems. In M. Ajmone Marsan, editor, *Application and Theory of Petri Nets 1993*, volume 691 of *Lecture Notes in Computer Science*, pages 1–16. Springer-Verlag, Berlin, 1993.
 13. T. Finin, J. Weber, G. Wiederhold, and et. al. Specification of the KQML Agent-Communication Language , 1993.
 14. F. Flores and J.J. Ludlow. Doing and Speaking in the Office. In *Decision Support Systems: Issues and Challenges*, pages 95–118. Pergamon Press, New York, 1980.
 15. D. Georgakopoulos, M. Hornick, and A. Sheth. An Overview of Workflow Management: From Process Modeling to Workflow Automation Infrastructure. *Distributed and Parallel Databases*, 3:119–153, 1995.
 16. C. Hagen and G. Alonso. Beyond the black box: Event-based inter-process communication in process support systems (extended version). Technical report, ETH Zürich, July 1997. Technical Report No. 303. <http://www.inf.ethz.ch/departement/IS/iks/publications/files/ha98c.pdf>.
 17. S. Jablonski and C. Bussler. *Workflow Management: Modeling Concepts, Architecture, and Implementation*. International Thomson Computer Press, London, UK, 1996.
 18. N. Jennings and M. Wooldridge, editors. *Agent Technology : Foundations, Applications, and Markets*. Springer-Verlag, Berlin, 1998.
 19. A. LaMarca, W.K. Edwards, P. Dourish, J. Lamping, I. Smith, and J. Thornton. Taking the Work Out of Workflow: Mechanisms for Document-Centered Collaboration. In *Proceedings of the Sixth European Conference on Computer-Supported Cooperative Work (ECSCW'99)*, 1999.

20. P. Lawrence, editor. *Workflow Handbook 1997, Workflow Management Coalition*. John Wiley and Sons, New York, 1997.
21. M. zur Mühlen. Evaluation of workflow management systems using meta models. In *Proceedings of the 32nd Hawaii International Conference on System Sciences - HICSS'99*, pages 1–11, 1999. <http://www.computer.org/proceedings/Hiccs2/0001/00010198Babs.htm>.
22. P. Muth, D. Wodtke, J. Weissenfels, A. Dittrich, and G. Weikum. From Centralized Workflow Specification to Distributed Workflow Execution. *Journal of Intelligent Information Systems*, 10(2), 1998.
23. K. Palacz and D.C. Marinescu. An agent-based workflow management system. In *Proc. AAAI Spring Symposium Workshop "Bringing Knowledge to Business Processes"*, Stanford University, CA, March 1999. <http://bond.cs.purdue.edu/docs/papers/awfms.ps>.
24. M. Reichert and P. Dadam. ADEPTflex: Supporting Dynamic Changes of Workflow without Loosing Control. *Journal of Intelligent Information Systems*, 10(2):93–129, 1998.
25. W. Reisig and G. Rozenberg, editors. *Lectures on Petri Nets I: Basic Models*, volume 1491 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, 1998.
26. W. Reisig and G. Rozenberg, editors. *Lectures on Petri Nets II: Applications*, volume 1492 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, 1998.
27. J. Rumbaugh, I. Jacobson, and G. Booch. *The Unified Modeling Language Reference Manual*. Addison Wesley, Reading, MA, USA, 1998.
28. J.R. Searle. *Speech Acts*. Cambridge University Press, Cambridge, 1969.
29. A.P. Sheth, W.M.P. van der Aalst, and I.B. Arpinar. Processes Driving the Networked Economy: ProcessPortals, ProcessVortex, and Dynamically Trading Processes. *IEEE Concurrency*, 7(3):18–31, 1999.
30. K. Swenson. Collaborative planning: Empowering the user in a process environment. *Collaborative Computing*, 1(1), 1994. <ftp://ftp.ossi.com/pub/regatta/JournalCC.ps>.
31. Action Technologies. *ActionWorkflow Enterprise Series 3.0 User Guide*. Action Technologies, Inc., Alameda, 1996.
32. Purdue University. Bond. the distributed object multi-agent system. <http://bond.cs.purdue.edu>, 2000.
33. E. Verbeek and W.M.P. van der Aalst. Woflan Home Page, Eindhoven University of Technology, Eindhoven, The Netherlands. <http://www.win.tue.nl/~woflan>.
34. E.M. Verharen, F. Dignum, and S. Bos. Implementation of a cooperative agent architecture based on the language-action perspective. In *Intelligent Agents*, volume 1365 of *Lecture Notes in Artificial Intelligence*, pages 31–44. Springer-Verlag, Berlin, 1998.
35. T. Winograd. Special Issue on the Language Action Perspective - Introduction. *ACM Transactions on Office Information Systems*, 6(2):83–86, 1988.
36. T. Winograd and F. Flores. *Understanding Computers and Cognition: A New Foundation for Design*. Ablex, Norwood, 1986.
37. M. D. Zisman. Use of production systems for modeling asynchronous concurrent processes. *Pattern-Directed Inference Systems*, pages 53–68, 1978.