Workshop Proceedings

The Process Modelling Group

(formerly "Petri nets and pi calculus for business processes")

Technische Universiteit Eindhoven

6-8 June 2005

Table Of Contents

Table Of Figures

Workshop Proceedings   *Page 2 of 21*
The Process Modelling Group
Technische Universiteit Eindhoven, 6-8 June 2005   3 September 2005

About This Document

This section explains the purpose, structure and scope of this document.

The **Process Modelling Group** (formerly known as **"**Petri nets and pi calculus for business processes") had its first face to face meeting at the Technische Universiteit Eindhoven from 6 to 8 June 2005.

The event, arranged by the group organizers and hosted by Wil van der Aalst, was a workshop rather than a conference. As such the atmosphere was informal. Varied presentations were given: some based on previously written material, some as informal explanations of work currently in progress, and some simply being adhoc talks inspired by previous discussion during the event. In general, presentations were short, with emphasis being placed on the subsequent discussion amongst participants.

This document is the full report of proceedings. It includes:

- A brief summary of the event

- The full list of participants

- Rapporteur notes from each session

- Personal observations by participants after the event

- The new and extended process modelling challenges that arose directly from the event.

Further information is available online, including the original process modelling challenges and solutions, as well as the full text of the written presentations:

- The Web forum and online document repository for internal group use is www.smartgroups.com/groups/petri_and_pi. This site is restricted to group members only.

- The Web site www.process-modelling-group.org is for use by the general public interested in our work, and includes a subset of the information available at the internal site.

Summary Of Proceedings

Given below is a brief summary of proceedings at the workshop.

The Petri-Pi Group had its first face-to-face meeting at the Technische Universiteit Eindhoven from 6 to 8 June 2005. What follows is a brief description of the group and summary of the event, for the interest of the general business process modelling community.

The group was instigated by Robin Milner, originator of the pi-calculus. It is run in collaboration by Wil van der Aalst, Rob van Glabbeek, Keith Harrison-Broninski, Robin Milner and Roger Whitehead. The membership (current size 70) includes people working in academia, for software companies and as part of standards bodies. The aims of the group, as agreed at the meeting, are two-fold.

The first aim is to put existing theoretical work to use in enhancing the quality of commercial process software. As a mechanism for this, we will help to provide a formal underpinning for mainstream process modelling languages and notations (**Business Process Execution Language** *aka* **BPEL** for machine workflow, **Choreography Description Language** *aka* **CDL** for machine interaction, **Role Activity Diagrams** *aka* **RADs** for human-driven processes, and so on) via the application of established formal techniques (Petri nets, pi calculus, Z notation, and so on). We aim to develop viable semantics for mainstream process modelling, and in so doing may be able to contribute usefully to the development of practical techniques. Group members include active participants in the committees responsible for the standards concerned, so the necessary communication channels are already in place.

The second aim is to further our scientific understanding of process modelling, and share the insights gained with the business community. As a mechanism for this, we are building a library of *challenges*—process examples, described in plain English text. Each group member is encouraged to model the challenges using techniques in which they are expert. By so doing, we hope not only to discover the respective strengths and weaknesses of different process modelling methods, but to develop *universal patterns for process modelling*—a catalogue of approaches, each applicable to particular situations. The initial set of challenges and solutions has already given rise to further more complex variants, which suggests to us that the challenge mechanism is a promising heuristic.

A manifesto for the group is in preparation, which expands on the aims described above to provide a long-term basis for our work. Those interested in joining the group should contact Roger Whitehead (rgw@office-futures.com) or Keith Harrison-Broninski (khb@rolemodellers.com). Please note that membership is open not to organizations or committees, but to *individuals* with proven interest and expertise in process modelling, regardless of their professional affiliation.

Participants

Listed below are all participants at the workshop. Our apologies to anyone who has been accidentally omitted – there were several last minute additions and cancellations.

Wil van der Aalst

Jos Baeten

Alistair Barros

Marek Bednarczyk

Carsten Butz

Rik Eshuis

Rob van Glabbeek

Keith Harrison-Broninski

Kohei Honda

Nickolas Kavantzas

Manuel Mazzara

Robin Milner

Natalya Mulyar

Karsten Schmidt

Christian Stefansen

Eric Verbeek

Mark Voorhoeve

Nobuko Yoshida

Rapporteur Notes

This section of the report summarizes each session of the workshop in turn, based on informal notes taken by the organizers. The actual papers presented are not reproduced here, but are available online for internal group use at www.smartgroups.com/groups/petri_and_pi.

Session 1 (Monday 6 June 2005, PM)

Robin opened the meeting. It was decided that Rob will produce a report for Monday afternoon, Robin for Tuesday morning, Wil for Tuesday afternoon and Keith for Wednesday morning (until he leaves). The possible long term aims of this meeting, and the efforts surrounding it, could be:

- Making one model for business processes

- Striving towards cohesion of different approaches

- Projecting a scenario for business processes that applies 10 years from now

… to mention a few possibilities. This will be discussed on Tuesday. Then we went over the program that would unfold over the next couple of days, and subsequently there was a round of introductions:

- **Wil van der Aalst** will introduce himself in the first talk.

- **Jos Baeten** works in process algebra and is also interested in Petri nets; he is interested in the strength and weaknesses of these approaches for workflow modelling.

- **Marek Bednarczyk** is hoping to find ways of putting mobility into Petri nets.

- **Carsten Butz** worked with XML and JAVA technology and is looking for a standard for business process modelling, in the Pi-calculus corner.

- **Rik Eshuis** is interested in Petri nets, statecharts and the Pi-calculus.

- **Rob van Glabbeek** is interested in integration of useful formalisms through mutual behaviour respecting translations.

- **Nickolas Kavantzas** is one of the designers of BPEL though OASIS.

- **Manuel Mazzara** is interested in the relationship of BPEL and process algebra and learning different approaches.

- **Robin Milner** is the originator of CCS and the pi calculus.

- **Natalya Mulyar** works on patterns for information systems.

- **Karsten Schmidt** wrote a model checker based on Petri nets.

- **Christian Stefansen** studies languages for business processes.

- **Erik Verbeek** works on semantic process models and is interested in verification.

Note: Some workshop participants were absent from this initial session and hence not included in the list above: Alistair Barros, Keith Harrison-Broninski, Kohei Honda, Mark Voorhoeve, and Nobuko Yoshida.

We discussed the timetable for the day. Marek would do a talk based on examples right after Wil. Then there would we an opportunity for anyone to present simple examples making a point. Robin volunteered to present a pi-calculus specification of the complaint-handling workflow. Rob would say something about using process algebra to specify the travel agent.

Wil's talk was on "Patterns and Nets". The www site www.workflowpatterns.com defines 20 basic control-flow patterns. Similarly, data patterns and resource patterns have been described. Each of these patterns collections has been used to evaluate workflow and other BPM products. In addition, Wil and his team have developed language-specific patterns such as CPN (Colored Petri Net) patterns. It seems that, in addition to "benchmark examples", patterns are a good way to compare and truly understand the various languages in the "Petri, Pi, BPEL space". Although Wil is a clear advocate of Petri net, he emphasizes the weaknesses of nets, e.g., the fact that the firing rule is local (no cancellation, no OR-join, etc.) and domain independent (no specific concepts for case, task, resource, etc.).

A few points of discussion:

- Christian wondered if the workflow patterns were made from a net-based perspective.

- Wil provided strong arguments against this feeling, including frequency-of-use studies.

- Christian also observed small ambiguities in the patterns; according to Robin that only helps to explain differences between approaches.

- Robin stated that states in process algebra are just as concrete as in nets: they are syntactic subexpressions.

- Jos added that in process algebra observable states can be modelled by propositional signals.

Marek's talk introduced Petri nets with mobility features, somewhat similar to Renew of Valk, and very disciplined. The approach was illustrated by an airplane example. The approach is kept simple by omitting features like deletion of agents.

Robin's complaint handling case study is documented on paper.

Note: See the group online document repository www.smartgroups.com/groups/petri_and_pi

Rob proposed using process algebraic operators that are optimized for certain applications like cancellation in the travel agent case study. These operators should fit in a specific format that guarantees that they can be expressed in terms of standard operators. Still the application dependent operators cannot be missed, because they lead to simpler specifications.

Session 2 (Tuesday 7 June 2005, AM)

Tuesday morning addressed two main topics.

The first topic was Web services and the pi calculus, and laid out a formal substrate for the W3C committee's work on choreography. First, Kohei Honda introduced pi, together with a broad description of how types are helpful in classifying behaviours. He then introduced the form of statement underlying **CDL**, the language adoped by W3C to provide a global view of an interactive system (Nick Kavantzas, the design leader for CDL, was also at the meeting.) These statements resemble those used in security protocols. In a second talk Nobuko Yoshida continued the theme, showing how CDL, with its **global** view, can be faithfully represented in pi calculus -- which in contrast emphasizes the **local** contribution made by each component to a global behaviour. Nobuko also gave more detail about session types, which constraint the intended behaviours of components.

*Note: Kohei and Nobuko have provided a new challenge based on their work with pi calculus and CDL, a Buyer-Seller-Shipper example that requires mobility. For details, see **Appendix A: New Challenge** below.*

Workshop Proceedings   *Page 7 of 21*
The Process Modelling Group
Technische Universiteit Eindhoven, 6-8 June 2005   3 September 2005

Copyright © 2005 The Process Modeling Group   **www.bptrends.com**

The second topic was the integration of different models. First, Rob van Glabbeek recalled a translation (due to Trakhtenbrot) of Petri nets into CCS, in which each transition is modelled as a pi process; he suggested another translation, more amenable (with the help of universal algebra) to the treatment of data, where instead each place is represented by a CCS process. Can this translation extend to mobile nets and pi calculus?

In a second talk, Keith Harrison-Broninski considered (with many examples) the distinction between mechanistic and human-driven processes. This led to considerable discussion about where the boundary lies between them, and whether research can shift the boundary. He went on to discuss three levels of control in a human-driven process: **management control**, **executive control** and **strategic control**. The latter, at present, does not have a formal basis. By contrast, he argued that pi calculus is appropriate for handling executive control, and Petri nets for management control.

Together, all these talks raised the hope that our different models are complementary, not conflicting, and that our job is to discover how use them all in concert.

Session 3 (Tuesday 7 June 2005, PM)

The session on Tuesday afternoon included 4 presentations and general discussion on the scope of the meeting and the results that should follow from it.

First was Alistair Barros from SAP Research in Brisbane, who gave a talk on Service Interaction Patterns. This new set of patterns is inspired by the workflow patterns initiative www.workflowpatterns.com and has a Web site www.serviceinteraction.com listing the patterns using Flash animations. Service interaction patterns allow emerging web services design and implementation solutions to be benchmarked against abstracted forms of representative scenarios. For example, the collected patterns can be used to evaluate languages and platforms supporting contract-based service development and service composition such as WSDL, WS-BPEL, JBI, Indigo, etc. The patterns go beyond simple bilateral interactions. They cover multilateral, competing, atomic, causally related, and routed interactions, as found in long-running business processes. The proposed solutions cover issues related to the implementation of these patterns using established and emerging web services standards and development frameworks. The focus of the patterns is on identifying principles, abstractions, and generic techniques.

The second talk was given by Christian Stefansen. He showed that using the pi-calculus' predecessor, Calculus of Communicating Systems (CCS), it is possible to correctly model the four original travel agency examples put forward by Wil. The challenges describe four versions of a travel agency handling flight, hotel, and car reservations with increasing degrees of complexity (optional booking, cancellation, roll-back). The solutions are nice but seem exponential in the size of the number of parties. After the raw algebraic encodings a sneak preview of SMAWL is provided. SMAWL is a workflow control language based on CCS. SMAWL aims to express all the currently identified control flow patterns in a natural way, by providing a simple set of high-level constructs. Christian's presentation shows that examples are vital for comparing different approaches.

The third talk was given by Rik Eshuis. His talk was similar to the talk given by Christian Stefansen. However, instead of CCS Rik used statecharts. Statecharts are a visual formalism for modelling reactive systems. They have been adopted in industrial standards like UML, for example. Statecharts differ from Petri nets in two ways: they have structured concurrency (AND/OR hierarchy) as well as reactive behaviour (event triggers). Rik presented a STATEMATE statechart solution for the complaint handling and travel agency cases (STATEMATE statecharts are the "original" version of statecharts as proposed by David Harel and his team, and are supported in the STATEMATE tool set). Rik's solution clearly shows that Statecharts have some

nice features not present in process algebra and Petri nets. Using the notion of the "nesting of states" it becomes easy to cancel tokens in a given region (this is the reason the cancellation sets were added to a language like YAWL). Nevertheless, statecharts also have things which as less easy as illustrated by the examples, e.g., leaving a state does not allow for some activities after the cancellation.

The fourth talk was given by Keith Harrison-Broninski. He presented an impromptu solution to the Travel Agency challenge using a Role Activity Diagram (RAD), extending it to include a customer, retail agent and wholesale agent, and in so doing generalizing the solution to include any number of booking requests. Beforehand he emphasized the importance of the human dimension. There were various questions on the semantics of RAD notation, for instance on iteration (also known as "replicated part refinement") and the scoping of data inside Roles. Keith explained that his approach to formal semantics for RADs is based on the Z notation, and that it may be possible to provide equivalent semantics using other techniques but that he had not investigated this. Robin said that there were strong similarities between RADs and the pi calculus which suggested that an equivalent semantics based on pi may well be possible.

Note: Following Keith's presentation, Robin requested that Keith submit the generalized process as a new challenge for use by the group, extending the original Travel Agency challenge. See **Appendix B: Extended Challenges** *below.*

Finally, there was discussion on the aims of the meeting and the intended results of all of this in the longer term. Items mentioned include:

- The need for mapping from one language to another.

- Learn from standards that failed and why they failed.

- Verification is already possible today, let us do it!

- There is considerable work on verification (often based on Petri nets) in the context of BPEL.

- BPEL should be the focal point of our efforts.

- Examples are very important.

- It would be good to not only collect examples and solutions but also a list of (verification) tools available today.

Session 4 (Wednesday 8 June 2005, Up To 11am)

The first session on Wednesday morning included 3 presentations.

First was Mark Voorhoeve, who gave an informal talk, "A graphical way of combining components", describing an approach he is developing for the graphical representation of processes that aims to integrate pi calculus and Petri nets in a single means of representation. High-level automata ("components") are shown with ports and connecting channels in a manner similar to the pi calculus, inside an enclosing higher-level component. States are not shown explicitly, but exist nevertheless within the components, each of which may correspond to a Petri net. Each component may be successively refined into more and more detailed Petri nets and finally an algebraic expression (based perhaps on the pi calculus) that represents an automaton in full detail. Mark raised a number of outstanding issues with this representation, to do for instance with the representation of mobility - it is work in progress.

Next, Robin Milner gave a talk on his new theory of **bigraphs**: "A model for information systems with mobile placing and linking". This adds to the pi calculus the concept of locality, something that is particularly useful for modelling *pervasive computing* devices - wearable computing

devices that respond to the computing machinery in the fabric of a building, for example. The theory of bigraphs aims not only to model spatial activity but also to unite diverse techniques such as Petri nets, lambda calculus, pi calculus, mobile ambients, and so on. The dynamics in the theory draw from diverse sources, for example the firing rules of Petri nets and the reaction/rewriting rules of the pi and ambient calculi. Robin has had some success so far in recreating existing techniques using bigraphs, for example finite CCS, pi calculus and mobile ambients.

The final talk of the session was an informal presentation by Nick Kavantzas of "A BPEL Solution to the Travel Agency Challenge". Nick showed this in a graphical rather than a textual representation, which omitted for clarity the mechanics of process invocation. There was some discussion of aspects of BPEL which are planned but as yet undefined, that may be important for real-world enactment of such processes - for example, **completion condition** and **parallel while**. Such features are unlikely to become part of the BPEL standard for about 12 months, and have not even been specified yet, so the group may be able to make a useful contribution to the BPEL committee discussions on these and other topics. Some differences between BPEL and CDL were also noted, for instance that "recovery" in CDL is more wide ranging than the equivalent failure handling mechanism "compensation" in BPEL. It was suggested that a BPEL "scope" may be equivalent in some sense to a bigraph "node".

As with previous sessions, the emphasis in all these talks was on integration of approaches: determining the modeling tasks that are most suited to each technique, and understanding how different techniques can be used in conjunction to complete the real-world process picture.

Concluding Session (Wednesday 8 June 2005, 11am to noon)

The workshop was a success with clear plans for the future. At the end the "work" was distributed as follows, although Roger was not there, and Keith and some others had to leave just before the end. Despite this the main agreements as discussed were:

(1) **Record of meeting (Keith Harrison-Broninski)**

Keith will take the lead in this. He will invite all participants to send: (1) a title and short abstract of their talk (max 200 words) and (2) general observations regarding the workshop (new insights, etc., max. 500 words). People will get 10 days for input and Keith will use this to make a report of the meeting. The goal is a factual recording of the event for those that were not there and also for future reference. The report should be first sent to the organizers, then to all attendants, and then to the whole SmartGroups list. The whole thing should be done in 1 month.

Note: This document is the result of the work described above.

(2) **Manifesto (Robin Milner)**

Robin will take the lead in writing a Manifesto discussing the aims of our efforts and the methods presented. The scope is on the interplay between formal methods (e.g. Petri and Pi) and languages from industry (e.g. BPEL). We will all forward some ideas. The Manifesto will refer to: (a) Examples and (b) Tools. Both will be a dynamic annex of the Manifesto.

*Note: The manifesto is still under discussion amongst the organizers. However, one initial conclusion that has been drawn is to rename the group, from the original but formalism-specific **Petri nets and pi calculus for business processes** to the more agnostic **The Process Modelling Group**.*

(3) **Examples (Wil van der Aalst)**

Wil will take the lead in getting more examples at the same level and editing these together with the solutions. The goal is to add same examples to the travel agency and the complaints handling, e.g., the examples in CDL and RAD. Marek Bednarczyk will assist in this.

(4) **Web site (Rob van Glabbeek)**

Rob will take the lead in building a website. Christian Stefansen will provide technical assistance and register "PetriPi.org".

Note: The Web forum and online document repository for internal group use only is www.smartgroups.com/groups/petri_and_pi. *It has since been decided to use the URL* www.process-modelling-group.org *in place of* www.petripi.org *for the public-facing Web site. Keith has registered the domains* **process-modelling-group.org**, **process-modelling-group.info** *and* **process-modelling-group.com**, *and set Web forwarding on all these to* www.petripi.org *(with URL hiding).*

(5) **Links to industry (Roger Whitehead)**

It seems that one of the problems is to get industry involved and make them aware of the strong tools and results that already exist. The hope is that Roger can play a role here as a liaison.

Note: Subsequently to the workshop, Roger requested a brief summary of the proceedings for reference in his column on the BPMG (Business Process Management Group) Web site www.bpmg.org. *Keith prepared this on 13 June 2005, all organizers had agreed it by 19 June 2005, and Roger then posted it online at* http://www.office-futures.com/PiPetri_report1.htm *for reference in his column. Keith also proposed that the summary be published on other relevant industry Web sites, and has since arranged with Paul Harmon for the summary to appear also on the industry Web site* www.bptrends.com.

Observations By Participants

This section of the report includes personal observations by participants of the workshop made after the event, edited only for formatting purposes. Hence the comments are not to be taken as representing the formal conclusions of the group, but solely as the opinions of the individual concerned.

Wil van der Aalst

A very nice meeting.

Common examples are a good way of comparing and truly understanding different languages.

BPEL seems to be there to stay so we should both accept this in the short term and propose more elegant languages for the long term.

It is nice to split the domain into mechanistic/human - interaction/process.

In the group we should think of means for rewarding people that are active and constructive (e.g., providing solutions for examples).

Rik Eshuis

The main focus of the workshop was to compare and exchange visual (Petri nets, statecharts, RAD) and textual (Pi calculus, CCS) techniques for modelling business processes. Most effort was spent in trying to understand each others solutions, rather than in analyzing them critically. My main interest in the workshop was to see the use of pi calculus in business process modelling. Having seen and heard several talks on this topic, and after having read some additional papers on pi, I have several observations and questions left, all regarding the use of pi calculus in business process modelling.

1) *Modelling*. The typical way to model a business process in pi calculus is to model the actors/resources involved and their interaction. The typical way to model a business process in Petri nets (and statecharts) is to model execution of a single case. Resources are typically not considered at this level (control flow only). This appears to me to be a fundamental difference. Perhaps session types can provide a bridge between these two views.

2) *Mobility vs evolution*. In the position papers, some people proposed using mobility for modelling process evolution. Process evolution is changing an active process definition. This is a very tough problem in workflow management since running cases can be half way the definition and the changes might result in a deadlock for those cases. Consequently, evolution cannot be allowed always (see e.g. the ADEPTflex system) I would therefore expect that mobility in pi calculus introduces similar problems, but this does not appear to be an issue (or perhaps this is where session types play a role).

   Moreover, the presented mobility examples do not address evolution, but rather flexibility of resources and case items, by passing their names along channels. In fact, in the examples the structure of the process does not seem to be rearranged. So I'd like to see an example pi process in which the control flow (so ordering in time) is really changed. If this is not possible, then it should become clearer what kinds of process evolution pi calculus supports.

3) *Expressiveness*. I had the feeling that most mobility examples simply involved parameter passing. This could also be done in say colored Petri nets. So the question is whether the full expressive power of the pi calculus is needed for business process modelling, or

whether a much less expressive subset with restricted name passing suffices. The same remark applies to using pi for giving a formal semantics to RADs: is the full-blown pi calculus needed for this, or does a stripped down version suffice?

In sum, the main result of the workshop for me is that it left with me more questions than I had before, and this is a positive thing.

Keith Harrison-Broninski

This group was formed at a critical juncture.

Awareness of the need for process orientation in the enterprise has been increasing among business people since the early 1990's, as a result of business transformation techniques such as **Business Process Reengineering** (**BPR**), management techniques such as **Six Sigma**, and pressure from standards bodies such as **ISO**. However, IT support lagged behind business needs until very recently.

With the start of the 21st century, however, came the beginnings of recognition in the software industry that enterprise computing must be *founded* on process principles. This recognition arose from two sources:

1. Ever-increasing complexity—something felt very painfully by many organizations when dealing with the Millennium Bug, and that is becoming more and more of a problem as competitive pressure due to globalization drives up the rate at which business practices in commercial organizations need to change;

2. The emergence of new Internet technologies—XML and Web services, for example—that made it possible to *scale up* process support to manage activities that span the entire enterprise.

Industry experts spread the word that enterprise-scale process support software was not only necessary but also now possible—and software vendors rushed to fill a vast new market. However, in the rush, the warnings of the experts that such software must be founded on sound theoretical principles were largely overlooked; incumbent vendors simply extended and re-badged legacy workflow and integration tools, while many new vendors claimed a theoretical foundation for their software that was only partially (if at all) valid.

Now, in 2005, the **Business Process Management** (**BPM**) market is established, worth many billions of dollars annually—and the problems are starting to emerge:

• Managing large and diverse processes is becoming an equivalent problem to the old one, of managing large and diverse systems;

• The robustness and security issues associated with cross-organizational processes can no longer be dealt with in an ad-hoc fashion;

• The vast amount of collaborative, innovative, flexible and dynamic human activity that forms the main body of work actually carried out in most organizations is not handled well, or at all, by current process support systems.

It has become a matter of urgency for the IT industry to get its house in order. This group may offer a way forward, since it brings together experts from across the board, with no commercial agenda to achieve or particular language standards to promote—and our three foci (machine workflow, machine interaction and human-driven processes) align directly with the 3 major business needs described above.

We are all in agreement that the way forward is both a practical and a scientific one—to develop a library of *process challenges*, and in solving them via different techniques determine and

promote the most appropriate manner to describe business processes. It is now the responsibility of all group members to contribute to this effort—solve challenges, analyse the solutions of others, work together to agree on a consistent way forward for process modelling in general, and spread the word to industry about our conclusions.

In particular, we must avoid the danger of becoming an academic debating group that has no relevance to the commercial world, by working closely with the existing efforts in industry to define and establish standards such as BPEL and CDL, and helping to establish any new standards that become necessary (e.g., for human-driven processes). If we are to achieve our aims, we must *engage with industry*—something that may be the biggest "process challenge" of all.

Kohei Honda

In terms of integration of applications (which may be among the concerns of BPM), web service would be a very good infrastructure. This may mean interaction use cases can be of increasing significance in future. If there are further scenarios in the integration setting, that would be a good thing to look at.

*Note: Kohei and Nobuko have provided such a use case as a challenge for use by the group. See **Appendix A: New Challenge** below.*

Robin Milner

The meeting worked even better than I had hoped. The main things were two.

First, we learned how many points of view there are, and how each deserves its place in some large consistent picture. For example, how can senior managers' views of a process be integrated with the meaning of the software that (partially) implements it?

Second, linked to the first, we met people who represent these points of view - not just idly, but in their professional lives.

So the most important outcome has been the initiation of a broadly based community working on a broadly ranging topic with immense challenges. What is vital now is to establish how this community will work. So the new Web site, and the activity it enables, will be our real test of success.

Christian Stefansen

The foremost point to be made is that our modelling efforts should be example/pattern-driven. Undoubtedly, Wil van der Aalst's collection of patterns on workflowpatterns.com has done a lot in terms of consensus convergence and so did the Travel Agency examples at the workshop. Examples should have a clear, unambiguous description of the problem, but - as Robin Milner pointed out - we should not require formal descriptions [of the challenges], because we then loose sight of the fine implementation nuances that different models provide.

Formal methods historically are an all or nothing proposition: "use our limited language, and we can provide guarantees". We should depart from this philosophy, and rather say "use your full language, and we will provide varying degrees of guarantees based on how many of the language features were actually used in your code". In the business modelling domain, formal tools need to address industry languages to gain impact, and the user should be made aware of the verification/expressiveness trade-off by offer rather than by nuisance.

Currently, the group is trying to make sense of the modelling languages that are already out there. Our primary goal should be to get involved and have sway over the modelling languages

Workshop Proceedings   *Page 14 of 21*
The Process Modelling Group
Technische Universiteit Eindhoven, 6-8 June 2005   3 September 2005

earlier in the process so we get avoid messy standards being adopted. The group needs to make a name for itself through publications and commentaries as well as through actively seeking out opportunities in standards committees. Our first step toward this end is the website www.petripi.org.

Note: The Web forum and online document repository for internal group use is www.smartgroups.com/groups/petri_and_pi. The Web site www.process-modelling-group.org (rather than the URL referenced above) is for use by the general public interested in our work.

Appendix A: New Challenge

*Kohei and Nobuko have provided a new challenge based on their work with pi calculus and CDL (see **Rapporteur Notes From Session 2** above).*

Buyer–Seller–Shipper Challenge

The new challenge is a Buyer-Seller-Shipper process that requires what in pi calculus terminology is called "mobility" – the ability to dynamically reconfigure connections during enactment of the process.

CHALLENGE:

A buyer asks a quote of a product to a seller.

The seller replies with a quote.

If the quote is acceptable to the buyer, she sends her credit card details as well as his private channel at which she wishes to receive the product (which we assume to be a digital datum).

The seller sends the credit card detail to a credit checking agency, which either replies with clearance of the credit data, or the converse. If the former (clearance) is the case, the seller sends an acknowledgement message to the buyer, and sends the product id and the buyer's private channel to the shipper, which delivers the product to that private channel, receives an ack, then reports that to the seller. On the other hand, if the credit data is not cleared, then the seller sends a refusal message to the buyer.

Notes On The New Challenge (Kohei Honda)

The use of a private channel may be understood as:

- Use of an existing private channel, known only to the buyer

- Creation of a new channel and its use

- Creation of a random number and its use.

It is a good idea to be able to model any one of these as formal descriptions.

While credit card data etc. are assumed to be confidentially delivered, this aspect does not have to be modelled. Modelling this aspect tractably is however an interesting challenge.

Other important variations:

(1) If Buyer does not like the quote it receives, it will ask another one. This can be either repeated arbitrarily, until some timeout fires.

(2) In the same vein, one may add a local timeout at each step when one party expects a communication of the other.

(3) Assume the same instances of interactions can occur concurrently with different buyers. Give a simple description in which the structure of interaction is clearly depicted regardless of involvement of one, two, three, or any number of buyers in the whole interaction.

There are many variants and refinements which arise naturally. Here only those are listed that are simple to comprehend as well as being typical.

Workshop Proceedings   *Page 16 of 21*
The Process Modelling Group
Technische Universiteit Eindhoven, 6-8 June 2005   3 September 2005

Appendix B: Extended Challenges

*Following Keith's impromptu solution using a Role Activity Diagram to the Travel Agency challenge (see **Rapporteur Notes From Session 3** above), extending and generalizing it in the process, Robin noted that this was an interesting outcome of the workshop and asked that Keith upload some new challenges to the Petri-Pi Web site for use by the group. This section describes the new challenges, and also shows corresponding solutions as Role Activity Diagrams in order to help clarify the textual explanations.*

Flexible Travel Agency Challenge

The first new challenge is a variant on the original Travel Agency challenge supplied by Wil van der Aalst. It arose as a result of doing an impromptu Role Activity Diagram solution to the Travel Agency challenge during the Eindhoven workshop. As so often in formal problems, the simplest solution was by making it more generic - allowing for n requests rather than just 3 (Flight, Hotel, Car) and adjusting the success criterion accordingly.

CHALLENGE:

This is a variant on the Travel Agency challenge, that makes it generic by allowing for n requests rather than just 3 (Flight, Hotel, Car) and adjusting the success criterion accordingly.

*In the variant challenge, it is possible for the customer to request **any number of items** from the Travel Agency, specifying a minimum number necessary. If this minimum number of items cannot be booked, the entire request is abandoned. Otherwise, if the minimum number are booked, payment is automatically made.*
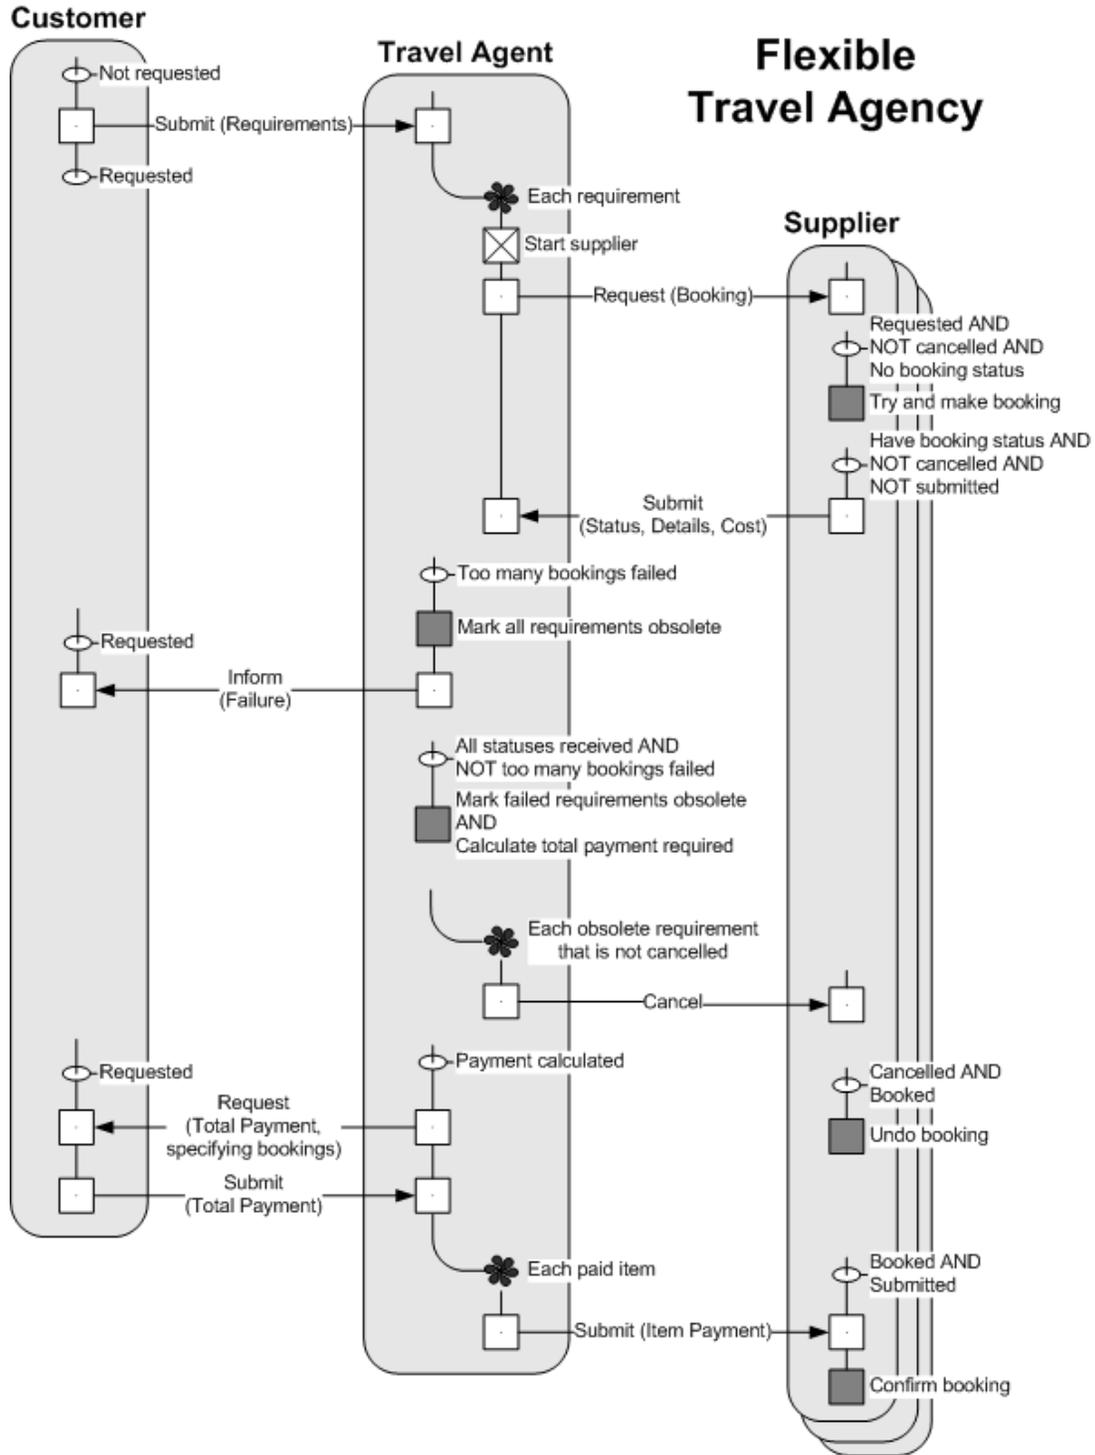
Customer

Travel Agent

Flexible
Travel Agency

- Not requested

Submit (Requirements)

- Requested

Each requirement

Start supplier

Request (Booking)

Supplier

Requested AND
- NOT cancelled AND
No booking status

Try and make booking

Have booking status AND
- NOT cancelled AND
NOT submitted

Submit
(Status, Details, Cost)

- Too many bookings failed

Mark all requirements obsolete

- Requested

Inform
(Failure)

All statuses received AND
- NOT too many bookings failed

Mark failed requirements obsolete
AND
Calculate total payment required

Each obsolete requirement
that is not cancelled

Cancel

- Requested

- Payment calculated

Cancelled AND
Booked

Request
(Total Payment,
specifying bookings)

Undo booking

Submit
(Total Payment)

Each paid item

Booked AND
Submitted

Submit (Item Payment)

Confirm booking

Figure 1: RAD Solution to Flexible Travel Agency Challenge

High Street Travel Agency Challenge

The second new challenge is a response to Robin's additional request that Keith supply challenges illustrating problems of process evolution. The second new challenge takes only an initial step in this direction - it is a variant of **Flexible Travel Agency** that attempts to model more closely a typical real-world Travel Agent. Essentially, a Travel Agent may deal with many Customers and many Suppliers at once, acting as a broker to co-ordinate their various dealings.

CHALLENGE:

This is a variant on the Flexible Travel Agency challenge, which attempts to model more closely a typical high street Travel Agent.

The Travel Agent may deal with many Customers at once. Each Customer registers with the Travel Agent, after which they may submit new requirements to the Travel Agent whenever they wish.

Each requirement is assigned by the Travel Agent to an appropriate Supplier to handle, with new Suppliers being contacted when necessary. Each Supplier may handle requirements from many Customers. Each Customer's requirements may be assigned to a variety of different Suppliers.

Each Supplier attempts to satisfy requirements in any order that they choose. When a Supplier manages to make a provisional booking for a requirement, it is sent directly to the Customer (rather than wasting time by going via the Travel Agent). The Customer may reject the booking by telling the Supplier, who will then cancel it. Alternatively, the Customer may accept the booking by providing the Supplier with payment details - the Supplier will then finalize the booking and pay commission to the Travel Agent.

Once the Travel Agent receives their commission, they (not the Supplier) will inform the Customer that the booking is confirmed. This is important. The Travel Agent confirming to the Customer means that the Customer gains some form of guarantee based on the Travel Agent's membership of travel industry organizations such as ABTA - a Supplier's confirmation would not provide such a guarantee.

Finally, each provisional booking has a time period (different in every case) during which it will be retained by the Supplier. If this period elapses without rejection or acceptance by the Customer, the Supplier is free to cancel it.
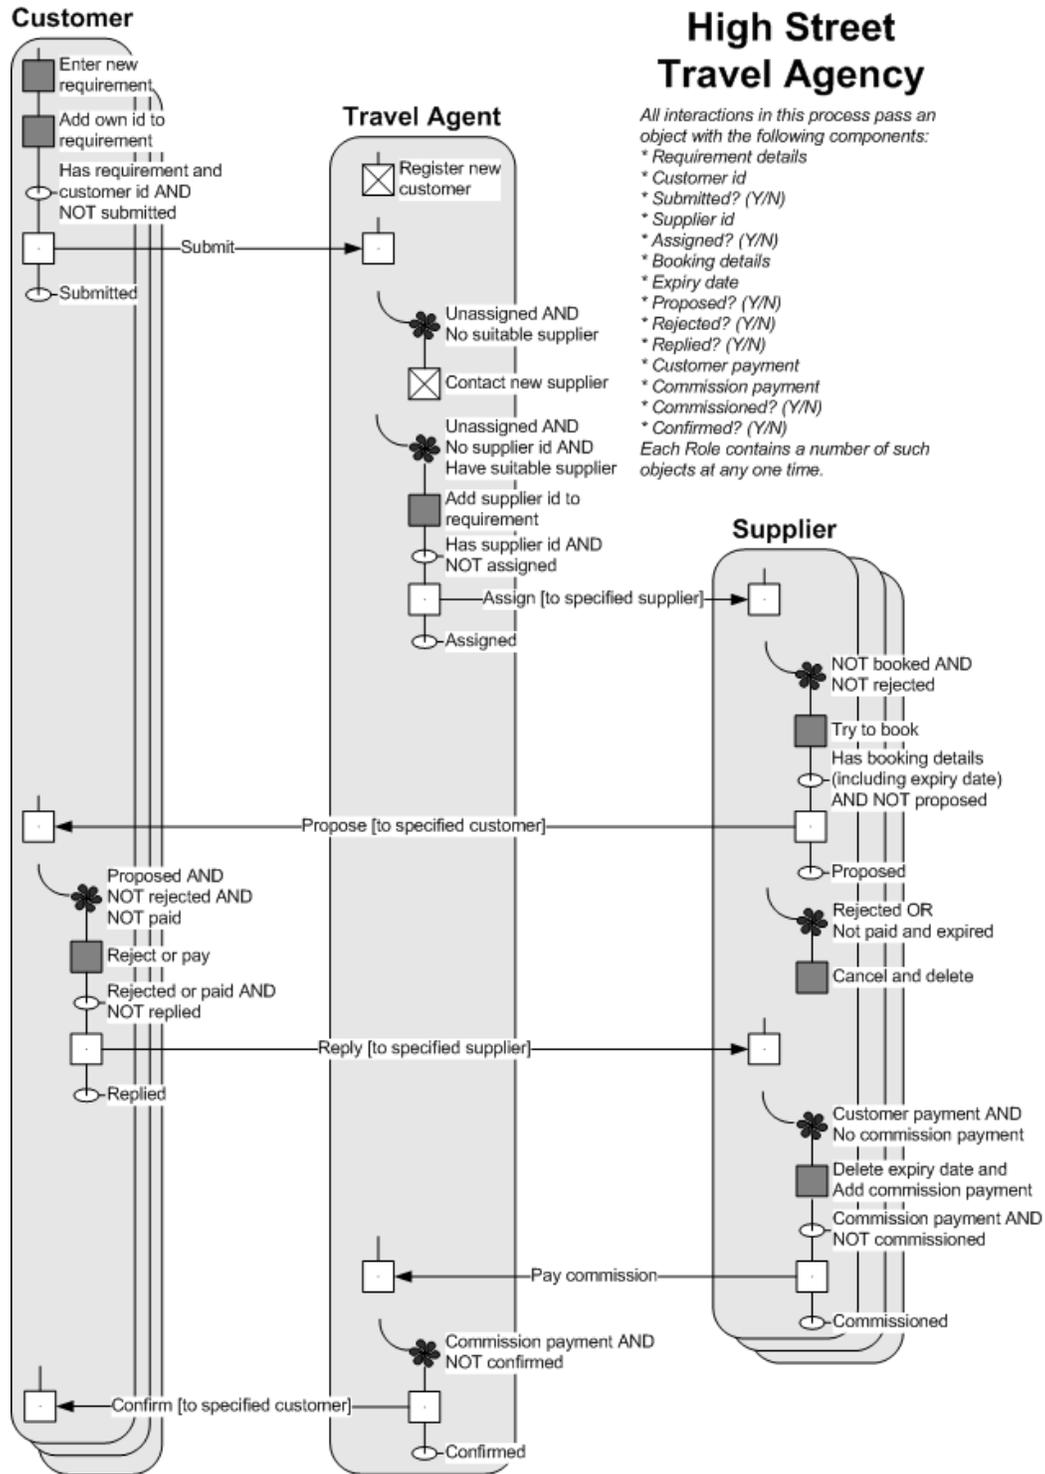
Figure 2: RAD Solution to High Street Travel Agency Challenge

Workshop Proceedings   *Page 20 of 21*
The Process Modelling Group
Technische Universiteit Eindhoven, 6-8 June 2005   3 September 2005

Notes On The Extended Challenges (Keith Harrison-Broninski)

The second new challenge takes the Travel Agency closer to real life - but still not close enough. It does not deal with information provision to the Customer, exception situations, the interaction between different bookings, successive refinement of bookings, payment options/problems, and so on. All these issues and more are a normal part of such processes - even when mechanized, say by a Web site.

Moreover, if we consider a human Travel Agent, the process becomes even more complex, since we are then in the domain of "human-driven processes", where the activities may go off in unforeseen directions. Consider a very normal situation: an agent who books business trips on behalf of a large company. All sorts of additional process issues arise, since if the agent wishes to prove their value and retain the company's business, they must effectively embed themselves into the internal business processes of the company concerned, taking note of considerations which differ from trip to trip and even making proactive suggestions based on their knowledge of the company's working practices.

However, we should take things one stage at a time. It is sensible to wait until we have some more solutions to the new challenges before extending them yet further. Once some we have some more solutions to the 2 new variants, we can make further extensions of the Travel Agency challenge which move yet closer to the real world.

Workshop Proceedings   *Page 21 of 21*
The Process Modelling Group
Technische Universiteit Eindhoven, 6-8 June 2005   3 September 2005