

# Expressing Workflow Patterns in CCS

Christian Stefansen

Department of Computer Science, University of Copenhagen (DIKU)  
Universitetsparken 1, DK-2100 Copenhagen Ø, Denmark

**Abstract.** There is an ongoing debate in the workflow community about the relative merits of Petri nets and  $\pi$ -calculus for workflow modeling. Recently, van der Aalst presented some challenges to model workflow in  $\pi$ -calculus. This paper responds to those challenges by showing how to code the 20 most commonplace workflow patterns in CCS (a subset of  $\pi$ -calculus), and describes two new workflow patterns that were identified in the process. The applicability of  $\pi$ -calculus to the workflow modeling domain is briefly discussed and a new *overlay* operator is discussed with applications to workflow descriptions.

## 1 Introduction

A prolonged debate has been going on about the relative merits of Petri nets [12] and  $\pi$ -calculus [11] for workflow modeling [21, 17]. Nonetheless, very little of this debate has been published or put in writing elsewhere. Recently van der Aalst presented seven challenges for those arguing in favor of  $\pi$ -calculus as a formal foundation for workflow systems [17]. This paper focuses on one of these challenges, namely:

**Challenge 6** Let the people that advocate Pi calculus exactly show how existing patterns can be modeled in terms of Pi calculus.

Whatever one's beliefs, a concrete encoding of the 20 workflow patterns will provide a basis for a more informed debate and make it easier for implementers to make a choice of formal foundation.

It is important to notice, however, that these workflow patterns are abstract patterns that do not make any assumptions about the data flow. This separation of concerns, besides being an important design principle, allows a clear and focused comparison without too much clutter. It also means that one can plug in one's data-flow language of choice at any later point. The analysis made here holds *regardless* of the data-flow language chosen, not just for one specific data-flow language, and so in effect the analysis is stronger and more general because of this abstraction<sup>1</sup>.

The following patterns may look a bit daunting when first presented, but one should remember that they are not what the workflow user/programmer should see any more than object calculus terms are what a Java programmer or UML user sees.

---

<sup>1</sup> Afterwards, a similar study could be made with regard to data-flow requirements, and of course any realistic system should handle both control flow and data flow.

## 1.1 Why Consider $\pi$ -calculus?

The  $\pi$ -calculus and its predecessor CCS (Calculus of Communicating Systems) [10] have strong mathematical foundations and have been widely studied for years. This has led to a plethora of deep theoretical results (see [15] for an overview) as well as applications in programming languages (e.g. PICT [13]), protocol verification, software model checking (e.g. Zing [1, 20]), hardware design languages<sup>2</sup>, and several other areas.

Workflow description languages, it would seem, have a lot in common with these areas. First, workflows can be thought of as parallel processes. Second, workflows often defy the block structure found in conventional programming. Third, the prospect of doing formal verification on workflows is very relevant and using a process calculus makes algebraic reasoning and algebraic transformation immediately possible due to the large body of research already present. Lastly, although they are somewhat bare in their style, CCS and  $\pi$ -calculus enforce a very strong separation of process and application logic that seems appropriate for workflow modeling.

On the other hand process calculi are highly theoretical constructs that require a great deal of expert knowledge. Workflow description systems should be accessible to anyone possessing a knowledge of the workflow domain being modeled, and so a significant challenge lies in bridging this gap. In particular providing understandable graphical tools and user-friendly abstractions constitutes a pivotal challenge if  $\pi$ -calculus-based systems are to succeed.

In the same vein a fair concern to raise is whether the level of formalism in  $\pi$ -calculus is necessary considering how graphical systems today are already difficult to understand both for users and programmers. First, using  $\pi$ -calculus does not preclude a high-level graphical representation, and clearly more research into this area is warranted. Second, a rigorous foundation remains necessary, because we strive to understand systems *before* we implement them. A well-tested, well-understood foundation will make implementation and subsequent adaptation easier and more routine.

A major strength of the  $\pi$ -calculus is its ability to express passing of channels between nodes and the ability to pass processes (in the higher-order  $\pi$ -calculus, which can be translated down to regular  $\pi$ -calculus). This far, however, only few examples exist that put this capability to use in the workflow domain. It may turn out, though, that channel-passing (and possibly locality) will become highly relevant when addressing service infrastructure and actual implementation of workflow tasks. By choosing  $\pi$ -calculus we stand a better chance of finding a unified foundation for business systems programming.

In terms of practical workflow management  $\pi$ -calculus promises seamless integration between workflow systems and existing verification tools (model checking [3, 8], behavior types/conformance checking [14, 9] etc.). For the end-user this means access to verification tools that will make writing and adapting workflows faster and less error-prone.

In our view it is important that future business systems are built not only to carry out a specific task, but in a way that makes them amenable to automatic source code manipulation and formal verification, and to attain this goal it is desirable to have a small and rigorously defined core model.

---

<sup>2</sup> In a sense, a microprocessor is one huge workflow system.

What we are seeking is not a formalism to use as it is, but the formalism that will provide the most suitable underpinning for future work.

## 1.2 Contributions

This paper demonstrates how CCS can be used as a foundation for workflow modeling. More specifically, the contributions are:

1. A CCS encoding of the 20 patterns in Table 1 as suggested in Challenge 6.
2. An brief analysis of the 20 patterns, and a description and encoding of two new patterns: (8a) *N-out-of-M Merge* and (16a) *Deferred Multiple Choice*.
3. A presentation of a speculative *overlay* operator and its use for certain workflow scenarios.
4. A short discussion of using CCS vs.  $\pi$ -calculus for workflow description languages.

The challenges suggest using  $\pi$ -calculus. It is our opinion that the foundation should be kept as simple as possible, and this research has in fact shown that CCS is sufficient for the purposes of the current workflow patterns.

## 1.3 Outline

In Section 2 we introduce the particular variant of CCS used throughout the paper. Section 3 offers a brief analysis of the workflow patterns as well as an encoding of the 20 common workflow pattern and the two newly identified ones. Section 3 also introduces the *overlay* operator. Section 4 discusses future directions, and Section 5 puts the paper in the context of related work. Section 6 contains a few concluding remarks.

## 2 Modeling Workflows in CCS

In this paper we will use the CCS syntax defined by the following BNF:

$$P ::= \mathbf{0} \mid \tau.P \mid a?x.P \mid a!x.P \mid P + P \mid (P \mid P) \mid \mathbf{new} \ a \ P \mid a?*x.P$$

$\mathbf{0}$  is the empty process,  $\tau$  is some unobservable transition,  $a!x$  and  $a?$  try to send and receive on channel  $a$ ,  $+$  is choice,  $\mid$  runs processes in parallel,  $\mathbf{new} \ a \ P$  protects the channel  $a$  from communication outside  $P$ , and  $a?*x.P$  spawns the process  $P$  each time a message is received on  $a$ . The syntax allows unguarded choice but guards all replicated processes with an input prefix ( $a?*x$ ). For simplicity we often write  $a?x$  instead of  $a?x.\mathbf{0}$ , and we may omit the name  $x$  if it is irrelevant in the context. Capital letters, usually  $P$ ,  $Q$ ,  $R$ , denote processes, and small letters,  $a$ ,  $b$ ,  $c$ ,  $\dots$  denote activities. Internal messages are words in small letters, like *ok* and *go*.

The invocation of an activity  $a$  can be encoded as  $a!.a?$  denoting a simple request/response mechanism. More sophisticated protocols for monitoring activity progress can be added, but for this paper the statuses started/finished will suffice. Often if  $a$  is a workflow activity we will simply write  $a$  for  $a!.a?$ . This will make the workflow activities easier to separate out from the internal message passing.

The non-determinism inherent in CCS should not be considered a problem; rather it is a convenient and elegant method of abstracting from implementation details, application logic or user input. The expression  $\tau.P + \tau.Q$  could be abstraction over a data-dependent choice to be made by the system or the decision of a human being.

### 3 The Workflow Patterns

In this section we consider each of the 20 workflow patterns in turn, discuss their encoding in CCS, and present new patterns. To facilitate comparison the discussion is structured according to the taxonomy given in [18] (see Table 1), and the pattern descriptions are taken verbatim from [19] with minor clarifications in square brackets.

<p><b>Basic Control Patterns</b></p> <ul style="list-style-type: none"> <li>1 Sequence</li> <li>2 Parallel Split</li> <li>3 Synchronization</li> <li>4 Exclusive Choice</li> <li>5 Simple Merge</li> </ul> <p><b>Advanced Branching and Synchronization Patterns</b></p> <ul style="list-style-type: none"> <li>6 Multiple Choice</li> <li>7 Synchronizing Merge</li> <li>8 Multiple Merge</li> <li>8a N-out-of-M Merge (<i>new</i>)</li> <li>9 Discriminator</li> <li>9a N-out-of-M Join</li> </ul>	<p><b>Patterns Involving Multiple Instances</b></p> <ul style="list-style-type: none"> <li>12 MI without synchronization</li> <li>13 MI with a priori known design time knowledge</li> <li>14 MI with a priori known runtime knowledge</li> <li>15 MI with no a priori runtime knowledge</li> </ul> <p><b>Structural Patterns</b></p> <ul style="list-style-type: none"> <li>10 Arbitrary Cycles</li> <li>11 Implicit Termination</li> </ul> <p><b>State-Based Patterns</b></p> <ul style="list-style-type: none"> <li>16 Deferred Choice</li> <li>16a Deferred Multiple Choice (<i>new</i>)</li> <li>17 Interleaved Parallel Routing</li> <li>18 Milestone</li> </ul>	<p><b>Cancellation Patterns</b></p> <ul style="list-style-type: none"> <li>19 Cancel Activity</li> <li>20 Cancel Case</li> </ul>
--	--	--

**Table 1.** The 20 Workflow Patterns [19] and Two New Ones

A central part of any workflow formalism is that of splitting the control flow into several possible paths and later joining control flow from different places into fewer paths. This is addressed by the *Basic Control Patterns*, the *Advanced Branching and Synchronization Patterns*, and to a certain extent the *Patterns Involving Multiple Instances*, the *Deferred Choice* pattern, and the *Interleaved Parallel Routing* pattern.

A structured view of these patterns proves beneficial. Consider Table 2 for an overview of Split, Synchronize, and Merge patterns. Split patterns split one path into some  $m$  paths. The difference between the split patterns is whether they require control flow to enter only *one* path (choice), *some* paths (multiple choice) or *all* paths (parallel split).

The same trichotomy is useful when considering synchronization and merging. Synchronization waits for a number of signals (one, some or all) and then launches one continuation, whereas merge spawns a continuation for each received signal (at most one, at most some number or for all).

Two new patterns emerge from this exercise: (16a) *Deferred Multiple Choice* and (8a) *N-out-of-M Merge*. These are described and coded below.

More generally, one could give every join construct a predicate that dynamically decided when to continue. If equipped with some simple constructs for writing such predicates, the language becomes much more expressive than the patterns here, but for analysis purposes it is very important that the predicates remain within the boundaries of Presburger arithmetic or some other (preferably simpler) first-order theory in which all statements are decidable.

	<b>Split</b>		<b>Synchronize</b>		<b>Merge</b>	
<b>All</b>	Parallel Split	$1 \rightarrow m/m$	Synchronization	$m/m \rightarrow 1$	Multiple Merge	$m/m \rightarrow m$
<b>One</b>	Exclusive Choice	$1 \rightarrow 1/m$	-	-	Simple Merge	$1/m \rightarrow 1$
	Deferred Choice	$1 \rightarrow 1/m$	Discriminator	$1/m \rightarrow 1$	-	-
<b>Some</b>	Multiple Choice	$1 \rightarrow n/m$	Synch. Merge	$n/m \rightarrow 1$	Multiple Merge	$n/m \rightarrow n$
	Deferred	$1 \rightarrow n/m$	N-out-of-M Join	$n/m \rightarrow 1$	N-out-of-M	$n/m \rightarrow n$
	Multiple Choice				Merge	

**Table 2.** Split, Synchronize, and Merge Patterns.  $l \rightarrow n/m$  means “control flows from  $l$  path(s) to  $n$  out of  $m$  possible paths”.  $n/m \rightarrow l$  means “control flows from  $n$  of  $m$  possible paths into  $l$ ”.

In the following pattern encodings some internal message-passing is often used for synchronization purposes. Such internal messages, like *done*, *ok*, *start*, and *go*, should be concealed to outside expressions through use of the **new** operator. Except for the first example, we tacitly assume their existence to avoid cluttering the expressions.

**Pattern 1. Sequence** – *execute activities in sequence*. The first encoding that comes to mind is  $a.P$ , but unfortunately this does not do the job completely. It is desirable to be able to put two arbitrary processes after each other like  $P.Q$ . A simple solution is provided by Milner [10]: We require all processes to send on an agreed-upon channel, say *done!*, when they are done and the encoding of the sequence  $P.Q$  then becomes

$$\text{new } go \{ \{ go^o / done \} P \mid go?.Q \}$$

where  $\{ go^o / done \} P$  is an alpha-conversion that is handled statically on source code level. Henceforth we will use the simpler notation  $P_{go}$  to signify that  $P$  signals on channel  $go$  on completion and that all pre-existing free *oks* in  $P$  have been alpha converted. If the channel is not relevant, it will be omitted. □

**Pattern 2. Parallel Split** – *execute activities in parallel*.

$$P_1 \mid \dots \mid P_n \quad \square$$

**Pattern 3. Synchronization** – *synchronize two parallel threads of execution.*

$$P_{1,ok} \mid \cdots \mid P_{n,ok} \mid \underbrace{ok? \cdots ok?}_n .R \quad \square$$

**Pattern 4. Exclusive Choice** – *choose one execution path from many alternatives.*

$$\tau.P_1 + \cdots + \tau.P_n$$

The  $\tau$  transition prefix on each branch allows a data-dependent decision or active decision; that is, the system can decide upon a particular branch (and do away with all other branches) without immediately activating the activities of that branch (cf. (16) *Deferred Choice*). In future patterns involving choice the  $\tau$  prefix will be omitted unless needed.  $\square$

**Pattern 5. Simple Merge** – *merge two alternative execution paths.*

$$(P_{1,ok} + \cdots + P_{n,ok}) \mid ok?* .R$$

The pattern assumes (unnecessarily for our purpose) that none of the processes  $P_i$  are ever run in parallel. This assumption is expressed here by the use of  $+$ .  $\square$

**Pattern 6. Multiple Choice** – *choose several execution paths from many alternatives.* A simple encoding would be

$$(\tau.P_1 + \mathbf{0}) \mid \cdots \mid (\tau.P_n + \mathbf{0})$$

but this (a) does not enforce a minimum number of activities to be executed and (b) does not explicitly tell if the system is still waiting for a choice to be made or already decided to take the  $\mathbf{0}$  branch. Addressing (b) we get

$$(\tau.P_1 + lazy!) \mid \cdots \mid (\tau.P_n + lazy!)$$

where the system outside should then accept messages on channels  $ok$  and  $lazy$  appropriately. Addressing (a) amounts to requiring some number of  $P_i$ s to be started before the system is able to perform rendez-vous on  $lazy$ . To avoid cluttering up the expression, we do not do this here.  $\square$

**Pattern 7. Synchronizing Merge** – *merge many execution paths. Synchronize if many paths are taken. Simple merge if only one execution path is taken.*

$$(P_{1,ok} + ok!) \mid \cdots \mid (P_{n,ok} + ok!) \mid \underbrace{ok? \cdots ok?}_n .Q$$

All  $P_i$  perform an  $ok!$  when they are done so the synchronization mechanism guarding  $Q$  should wait for exactly  $n$  such messages.  $\square$

**Pattern 8. Multiple Merge** – *merge many execution paths without synchronizing.*

$$(P_{1,ok} + lazy!) \mid \cdots \mid (P_{n,ok} + lazy!) \mid ok?* .Q \mid lazy?*$$

Signals *ok!* and *lazy!* signify if an activity was executed or not. Notice that other processes may connect to the merge by using the named entry-point *ok*. If this is undesired in the context, a new *ok* can be added around the expression. Additionally, any occurrences of *ok* in *Q* should be removed by alpha conversion.  $\square$

**New Pattern 8a. N-out-of-M Merge** – merge many execution paths without synchronizing, but only execute subsequent activity the first *n* times. Remaining incoming branches are ignored.

$$(P_{1,ok} + lazy!) \mid \cdots \mid (P_{n,ok} + lazy!) \mid \underbrace{ok?.go! \mid \cdots \mid ok?.go!}_n \mid go?* . Q \mid lazy?*$$

The pattern is similar to *Multiple Merge* except that it is only capable of receiving *n* messages on *ok* and then it is done.  $\square$

**Pattern 9. Discriminator** – merge many execution paths without synchronizing. Execute the subsequent activity only once.

$$(P_{1,ok} + lazy!) \mid \cdots \mid (P_{n,ok} + lazy!) \mid ok?.(Q \mid lazy?* \mid ok?*)$$


Here the *Multiple Choice* pattern is used but other split patterns could be used too. The first signal on *ok* initiates *Q* and all remaining signal are consumed. Contrary to Petri nets, ignoring all future signals is not a problem because any loop around this construct would simply instantiate a new process (with new channel names) in each iteration.  $\square$

**Pattern 9a. N-out-of-M Join** – merge many execution paths. Perform partial synchronization and execute subsequent activity only once.

$$(P_{1,ok} + lazy!) \mid \cdots \mid (P_{n,ok} + lazy!) \mid \underbrace{ok? \dots ok?}_n . (Q \mid lazy?* \mid ok?*)$$

For the purpose of this demonstration, the pattern is combined with *Multiple Choice* pattern. Once *n* messages are received on *ok*, the process *Q* is activated and all remaining messages, whether *ok* or *lazy*, are discarded.  $\square$

**Pattern 10. Arbitrary Cycles** – execute workflow graph w/out any structural restriction on loops. Consider the following example from [19]:



Since activities *b* and *c* have multiple merge points that do not originate from the same split block, it is necessary to promote them to named services (or “functions”) by

using the replication operator – a lot like what one would do in any structured programming language. Here we name them  $go_b$  and  $go_c$  and the example becomes

$$a.go_b! \mid go_b?*.b.go_c! \mid (go_c?*.c.(go_b! + d.(go_c! + e))). \quad \square$$

**Pattern 11. Implicit Termination** – *terminate if there is nothing to be done.*

Implicit termination means detecting if all processes have come to a stable state and there are no pending activities or messages. In other words if the residual workflow process – disregarding subexpressions of the form  $a?*.P$  – is structurally congruent to  $\mathbf{0}$ . Subexpression of the form  $a?*.P$  can be thought of as functions; they cannot and should not be reduced. Any active invocations of such a function will manifest themselves as some reduced form of the body  $P$ . (However, if a function listens on a channel that cannot be reached by the active process expression, it may be garbage collected by the runtime system, but this is separate from detecting implicit termination.)

Explicit termination was modeled in the *Sequence* pattern where we adopted the convention that all processes signal on a pre-determined channel on completion. When the top-level expression wishes to send its signal, the process explicitly terminates.  $\square$

**Pattern 12. MI without synchronization** – *generate many instances of one activity without synchronizing them afterwards.*

$$P.loop! \mid loop?*. (create!.loop! + R) \mid create?*.Q$$

After process  $P$  finishes an arbitrary number of instances of  $Q$  are spawned and once no more  $Q$ s need to be started,  $R$  is executed. Execution of  $R$  does not wait for the completion of any of the instances of  $Q$ . The choice inside the loop can be an *Exclusive Choice* or a *Deferred Choice*.  $\square$

**Pattern 13. MI with a priori known design time knowledge** – *generate many instances of one activity when the number of instances is known at the design time (with synchronization).*

$$\underbrace{create! \dots create!}_n . \underbrace{ok? \dots ok?}_n . Q \mid create?*.P_{ok}$$

This pattern creates exactly  $n$  instances of  $P$  and waits for all of them to complete before passing control on to the process  $Q$ .  $\square$

**Pattern 14. MI with a priori known runtime knowledge** – *generate many instances of one activity when a number of instances can be determined at some point during the runtime (as in FOR loop but in parallel).* First a counter is needed to keep track of the number of instances that need to be synchronized. It can be coded (a) directly in CCS through cunning use of the `new` operator (cf. sect. 7.5 of [10]), (b) by adding a notion of simple datatypes such as integers to CCS or (c) by using channel-passing  $\pi$ -calculus style. For now we will just add a process *Counter* and not bother with its implementation details; the important part being that is it feasible in CCS. *Counter*



has the ability to receive messages on channel *inc* (increase) and to send message on channels *dec* (decrease) or *zero* (check if zero) depending on its state:

$$P.loop! \mid loop?*(create!.loop! + break!) \mid create?*!.inc!.Q_{ok} \\ \mid break?*(ok?.dec?.break! + zero?.R) \mid Counter$$

From  $P$  the process enters a loop where multiple instances of  $Q$  are created. Once creation is done, the process iterates the *break* loop until a message on *zero* is received, i.e. 0 instances remain. Notice that one can decide not to create any instances at all., if one desires a minimum number of processes to be spawned, the *break!* prefix should be guarded by unfolding the *create* loop an appropriate number of times.  $\square$

**Pattern 15. MI with no a priori runtime knowledge** – *generate many instances of one activity when a number of instances cannot be determined (as in WHILE loop but in parallel)*. This pattern is merely a simplified version of pattern 14. There is no longer a need to distinguish between the creation phase and the synchronization phase. Hence it collapses to:

$$P.loop! \mid loop?*(create!.loop! + ok?.dec?.loop! + zero?.R) \mid create?*!.inc!.Q_{ok} \\ \mid Counter$$

More advanced synchronization schemes can be plugged in at this place. Maybe only some  $n$  instances need to finish, maybe the completion condition is some predicate  $\rho$  over the data produced so far. In other words, the logic deciding this loop can be pushed up to a data-aware layer or handled through more complex join conditions in the process expression.  $\square$

**Pattern 16. Deferred Choice** – *execute one of the two alternatives threads. The choice which thread is to be executed should be implicit.*

$$P_1 + \dots + P_n \quad \text{where each subprocess is guarded.}$$

This is very similar to the *Exclusive Choice*. Here the choice is made exactly when a (non-silent) transition of either of the  $P_i$ s occurs, and hence we require all  $P_i$  to be guarded. In the *Exclusive Choice* we might have  $\tau.P + \tau.Q$  to signify that an abstract upper-layer would decide which branch to follow (i.e. what  $\tau$  transition to take).  $\square$

**New Pattern 16a. Deferred Multiple Choice** – *execute several of the many alternative threads. The choice of which threads are to be executed should be implicit.* The important question here is: how do we know when the users are done choosing the threads? Two approaches are possible: (a) we fix an  $n$  number of threads to be activated or (b) we set up a virtual activity that means “all desired threads have been started, remove remaining choices”. Option (a) requires all subprocesses to emit a signal on activation and consume  $n$  such signals before activating a *lazy?\** process to remove the remaining branches. Here we follow option (b):

$$(P_1 + lazy!) \mid \dots \mid (P_n + lazy!) \mid donechoosing.lazy?*$$

Firstly, this pattern cannot be obtained by combination of several *Deferred Choice Patterns*. For instance, the expression  $(A.(B+donechoosing)+B.(A+donechoosing)+donechoosing)$ , does not allow the second activity to be started before the first one is done.  $\square$

**Pattern 17. Interleaved Parallel Routing** – *execute two activities in random order, but not in parallel.*

$$\text{new } lock, unlock, locked, unlocked (lock!.P_1.unlock! | \dots | lock!.P_n.unlock! \\ | (unlocked? * .lock?.locked! | locked? * .unlock?.unlocked! | unlocked!))$$

This is generalized slightly from two to any finite number of interleavings. Each process  $P_i$  acquires the lock on activation and releases it upon termination.  $\square$

**Pattern 18. Milestone** – *enable an activity until a milestone is reached.* Assume that process  $Q$  can only be enabled after  $P$  has finished and only before  $R$  is started. This is done by the help of a small flag that can be changed using  $set!$  and  $clear!$  and tested using  $ison?$  and  $isoff?$ :

$$\text{Milestone}(ison, isoff, set, clear) = on? * .(ison!.off! + set?.on! + clear?.off!) | \\ off? * .(isoff!.off! + set?.on! + clear?.off!) | off!$$

Now the milestone can be set up as

$$P.set!.clearR | ison?* .Q | \text{Milestone}$$

where  $clearR$  denotes the modification of process  $R$  that signals on  $clear$  when it starts its first activity or makes its first explicit choice.  $\square$

**Pattern 19. Cancel Activity** – *cancel (disable) an enabled activity.* Suppose there are activities  $a$ ,  $b$ , and  $c$  being carried out in sequential order. Anytime during the execution of activity  $b$  the process can be cancelled. We break up activity  $b$  into  $b!.b?$  so it will be clear exactly when a cancellation can be accepted. Assume cancellation is obtained by signaling on channel  $cancel$ :

$$a.b!.(b?.c + cancel?) \quad \square$$

**Pattern 20. Cancel Case** – *cancel (disable) the process.* The example of this pattern given at [19] is an online travel reservation system. If the reservation for a plane ticket fails, then all pending reservations for flights, hotel, car rental, etc. in the itinerary must be cancelled immediately to avoid unnecessary work.

Intuitively, this corresponds to returning to some predefined state in the system after relinquishing all resources and information bound in the current context. In this respect the cancellation patterns look at lot like a program exceptions, and indeed the cancellation pattern can be coded in CCS by using a source code transformation.

Pattern 19 showed how to insert a cancellation point a one particular spot in a process. We now insert cancellation points for all states in the process  $a.b.c$  to obtain

$$a!.(a?.(b!.(b?.(c!.(c?.cancel?) + cancel?) + cancel?) + cancel?) + cancel?) + cancel?$$

Clearly this becomes very tedious and strongly suggests that a more abstract language would be useful.  $\square$

### 3.1 Towards a Workflow Description Language Based on CCS

As could be expected the workflow patterns coded in the previous sections clearly show that CCS is too low-level to be used directly as a workflow language, and that certain patterns are very cumbersome to write.

A natural idea is to design a language out of small building blocks of each of the patterns seen to far. The language SMAWL, presented in the related paper [16] seeks to reduce the amount of user-specified internal synchronization mechanisms and provide elegant constructs for the 22 workflow patterns while maintaining a strong link to CCS. The example in Figure 1 based on [19] shows what a workflow specification might look like in SMAWL along with an automatically generated graphical representation.

```

workflow Become a recording star =
  chooseone {
    ⇒ call (Work your way up)
    ⇒ call (Try to get lucky)
  };
  Make record;
  doall {
    ⇒ chooseone {
      ⇒ Develop as an artist
      ⇒ Develop bad habits
    }
    ⇒ Rehearse tour;
    ⇒ Do tour
  }
end
  
```

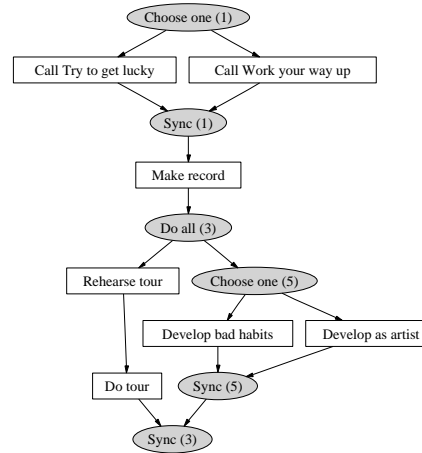



Fig. 1. How To Become a Recording Star (adapted from the *Recording Star* example [19])

### 3.2 Challenge 4 and the Overlay Operator

What remains is a tiny – somewhat speculative – improvement of the expressive power of CCS. Consider the Petri net in Figure 2. The Petri net exhibits a pattern that is cumbersome to express without a lot of internal message-passing because it defies the block structure of the workflow patterns.

The Petri net can be mimicked in CCS through use of function definitions and explicit message-passing, but more interestingly it can be expressed elegantly using the overlay operator & defined by the rules in Figure 3. The operator is somewhat inspired



**Fig. 2.** Challenge 4 [17]

by the  $\parallel$  operator found in CSP [7]. The operator is perhaps best explained by coding the example in Figure 2:

$$C4 = a.(b.c.d|e.f.g).h \& (c|e).f$$

Whenever  $C4$  wants to transition on a channel that is mentioned in both operands to the  $\&$  operator, both operands have to agree to this. If a transition is only mentioned in one of the operands, the  $\&$  operator has no effect on that transition. Here if  $a, e$  was performed first, the left process would be ready to do  $f$  but the right would not, since it would need  $c$  to complete first.

The operator is very interesting because it allows the overlaying of global business rules to all processes. E.g. one might have a rule that says “delivery shall occur before invoicing”. This rule would now simply be enforced on all workflows by overlaying it with the workflows. Also, the workflow designers do not have to describe all business rules because they can simply overlay the global rules to their workflows.

The operator can be built into the reduction rules of a workflow system, or it can be translated down to regular CCS, which, unfortunately, can generate an exponential blowup in expression size.

$$\frac{P \xrightarrow{\alpha} P' \quad \alpha \notin \text{ports}(Q)}{P \& Q \xrightarrow{\alpha} P' \& Q} \text{ (+sym.)} \qquad \frac{P \xrightarrow{\alpha} P' \quad Q \xrightarrow{\alpha} Q'}{P \& Q \xrightarrow{\alpha} P' \& Q'}$$

**Fig. 3.** Reduction Rules for the Overlaying Operator  $\&$

## 4 Discussion and Future Work

Challenges 4–7 were to (4) model the Petri net shown in Figure 2, (6) model existing patterns, (5) provide new challenges, and (7) suggest new patterns (especially using mobility).

We have modeled existing patterns and in doing so identified two new patterns. The two new patterns (*16a*) *Deferred Multiple Choice* and (*8a*) *N-out-of-M Merge* should be

expressed in Petri nets, and likewise it would be interesting to see if Petri nets can handle the overlaying operator more elegantly than CCS. An investigation into the practical applicability of the overlaying operator  $\&$  might prove very fruitful.

Also notice how we can express variants of existing patterns, e.g. we can combine deferred and non-deferred choices (it is not clear if this is possible using the current patterns).

$\pi$ -calculus mobility was not needed. Essentially, workflow systems deal with the flow of command/control and not as much the channels for exchanging messages, if one desires to shift focus to the flow of data—and in particular which channels are used for this—then maybe  $\pi$ -calculus is appropriate. As noted previously it would seem that the  $\pi$ -calculus notion of mobility would be convenient for modeling service infrastructure and implementation of actual workflow tasks.

When debating expressiveness an easy, but vacuous point is Turing completeness. Indeed, both Petri nets and CCS are Turing complete, but in their raw form they are unsuited for any non-trivial programming/modeling task. For domain-specific languages it would seem more reasonable to consider criteria such as (a) ontological fit (ie. domain closeness, convenience of expressing common idioms), (b) expressiveness, (c) amenability to formal analysis, (c) simplicity, and, for workflow languages, (d) graphical representation.

It remains extremely important to provide an intuitive graphical user interface and easily understandable high-level abstractions for the user, and this should be addressed in future work. For a demonstration of how to provide a more high-level syntax refer to the workflow language SMAWL [16].

To explore expressiveness properly, it would be interesting to put all the languages into a formal semantic framework. For workflows, the idea of representing workflows as sets of allowable traces CSP-style seems reasonable and will provide a common ground for comparison. First the patterns should be described in a more formal fashion and then the expressiveness of Petri nets, CCS, CSP, temporal logic and others can be compared using e.g. Felleisen characterization of expressiveness [5]<sup>3</sup>.

## 5 Related Work

Dong and Shensheng presented their encoding of some of the patterns in an unpublished paper in 2004 [4]. Their encodings differ a great deal from the ones presented here, but more fundamentally they use the channel-passing facility of  $\pi$ -calculus, whereas the encodings here do not and thus can be expressed in CCS. In our view the language should be as simple as possible, and as we have demonstrated, the full power of  $\pi$ -calculus is not needed. Also, staying within CCS makes the patterns more amenable to use in the tools current available. Furthermore the Join-calculus would seem interesting [6].

Numerous attempts to map Petri nets to process calculi exist. Also an effort to merge the best of both worlds exists (cf. Petri Box Calculus [2]).

---

<sup>3</sup> Quoting Felleisen [5]: . . . “more expressive” means that the translation of a program with occurrences of one of the constructs  $c_i$  to the smaller language requires a global reorganization of the entire program.

## 6 Conclusion

The purpose of this paper was to show that workflow patterns in fact can be expressed in CCS. We demonstrated how to code challenges 4 and 6 and identified new patterns.

The response to challenges 5 and 7 is necessarily this: The coding of the 20 patterns has provided no compelling reasons to use the channel-passing mechanism of  $\pi$ -calculus. We see no need for the notion of mobility found in  $\pi$ -calculus for workflow systems. Workflow systems abstract away from the actual channels of delivery. Maybe  $\pi$ -calculus is relevant when analyzing the interaction between agents and clients inside the workflow activities.

It remains open if the newly discovered patterns occur often enough in practical settings to warrant their incorporation with the rest. Real world examples to support them would be very valuable – especially for the overlaying operator.

## References

1. Tony Andrews, Shaz Qadeer, Sriram K. Rajamani, and Yichen Xie. Zing: Exploiting program structure for model checking concurrent software. In Philippa Gardner and Nobuko Yoshida, editors, *CONCUR*, volume 3170 of *Lecture Notes in Computer Science*, pages 1–15. Springer, 2004.
2. Eike Best, Raymond Devillers, and Maciej Koutny. *Petri Net Algebra*. Monographs on theoretical computer science. Springer, 2001.
3. E. M. Clarke, E. A. Emerson, and A. P. Sistla. Automatic verification of finite state concurrent system using temporal logic specifications: a practical approach. In *POPL '83: Proceedings of the 10th ACM SIGACT-SIGPLAN symposium on Principles of programming languages*, pages 117–126. ACM Press, 1983.
4. Yang Dong and Zhang Shensheng. Modeling workflow patterns with pi-calculus. Unpublished. Available from <http://www.workflow-research.de>.
5. Matthias Felleisen. On the expressive power of programming languages. *Science of Computer Programming*, 1990.
6. Cédric Fournet and Georges Gonthier. The reflexive chemical abstract machine and the join-calculus. In *Proceedings of the 23rd ACM Symposium on Principles of Programming Languages*, pages 372–385, St. Petersburg Beach, Florida, January 21-24 1996. ACM.
7. C.A.R. Hoare. *Communicating Sequential Processes*. International Series in Computer Science. Prentice-Hall, 1985.
8. Michael R. A. Huth and Mark Ryan. *Logic in computer science: modelling and reasoning about systems*. Cambridge University Press, 2000.
9. Naoki Kobayashi. Type systems for concurrent programs. In Bernhard K. Aichernig and T. S. E. Maibaum, editors, *10th Anniversary Colloquium of UNU/IIST*, volume 2757 of *Lecture Notes in Computer Science*, pages 439–453. Springer, 2002.
10. Robin Milner. *Communication and Concurrency*. International Series in Computer Science. Prentice-Hall, 1989.
11. Robin Milner. *Communicating and Mobile Systems: The  $\pi$ -calculus*. Cambridge University Press, 1999.
12. High-level Petri nets – concepts, definitions and graphical notation, final draft international standard ISO/IEC 15909, May 2002. Version 4.7.3.
13. Benjamin C. Pierce and David N. Turner. Pict: A programming language based on the pi-calculus. In G. Plotkin, C. Stirling, and M. Tofte, editors, *Proof, Language and Interaction: Essays in Honour of Robin Milner*. MIT Press, 2000.

14. Sriram K. Rajamani and Jakob Rehof. Conformance checking for models of asynchronous message passing software. In Ed Brinksma and Kim Guldstrand Larsen, editors, *CAV*, volume 2404 of *Lecture Notes in Computer Science*, pages 166–179. Springer, 2002.
15. D. Sangiorgi and D. Walker. *The  $\pi$ -calculus: a Theory of Mobile Processes*. Cambridge University Press, 2001.
16. Christian Stefansen. SMAWL: A SMALL Workflow Language based on CCS. In *Proceedings of the CAiSE Forum of the 17th Conference on Advanced Information Systems Engineering*, June 2005.
17. W.M.P. van der Aalst. Pi calculus versus Petri nets: Let us eat "humble pie" rather than further inflate the "Pi hype". Available from <http://tmitwww.tm.tue.nl/research/patterns/download/pi-hype.pdf>, 2004.
18. W.M.P. van der Aalst and A.H.M. ter Hofstede. Workflow patterns: On the expressive power of (petri-net-based) workflow languages. In K. Jensen, editor, *Proceedings of the Fourth Workshop on the Practical Use of Coloured Petri Nets and CPN Tools (CPN 2002)*, volume 560, Aarhus, Denmark, August 2002. DAIMI.
19. Workflow patterns. Available from <http://www.workflowpatterns.com>.
20. The Zing model checker. Available from <http://research.microsoft.com/zing/>.
21. Michael zur Muehlen. Workflow research. <http://www.workflow-research.de>.