Communication Abstractions for Distributed Business Processes

Lachlan Aldred¹, Wil M.P. van der Aalst^{1,2}, Marlon Dumas¹,
and Arthur H.M. ter Hofstede¹

¹ BPM Group, Queensland University of Technology, Australia
`{l.albred,m.dumas,a.terhofstede}@qut.edu.au`

² Department of Mathematics and Computer Science, Eindhoven University of
Technology, The Netherlands
`w.m.p.v.d.aalst@tue.nl`

Abstract. Languages for business process definition generally suffer from myopic approaches to capturing communication between distributed processes. Effective communication between processes requires: support for conversations involving interrelated interactions spread over time; ability to select and group messages based on their content, regardless of format and transport technology; and resolving contention between processes or tasks for common sets of messages. This paper presents a set of communication abstractions that provide a “glue” between the process layer and the middleware. The paper also reports on an implementation of these abstractions and an experimental evaluation.

Keywords: Business process integration, correlation patterns.

1 Introduction

At present, the definition of business processes that interact with one another in a distributed environment is hampered by a number of factors. Firstly, these processes are required to run on top of mainstream communication middleware which often does not support key features needed by applications in general [4,10,12,11], and process-oriented applications in particular. For instance message selectors in JMS¹ cannot filter messages based on their body. Secondly, there exist conceptual problems with state-of-the-art business process definition languages in regards to process-to-process communication abstractions. To receive message batches for example, Business Process Management Systems (BPMSs) force designers to incorporate dedicated code into the same scope as business process logic. Finally, apart from correlation mechanisms for routing messages to process instances, BPMSs accept messages sent to them without filtering, thus forcing unnecessary amounts of message selection code into the process definition. In summary, the distinction between process abstractions and communication abstractions is blurred in existing approaches.

¹ Java Message Service: java.sun.com/products/JMS accessed November 2006.

This paper motivates and defines a set of communication abstractions at a layer in-between the business process and the middleware, that simplify the definition of interactions between distributed processes. The proposed abstractions have been implemented on top of a communication API, namely JCoupling [5], that abstracts away from the underlying middleware and protocols. Using this implementation, we have conducted preliminary experiments to compare this approach with respect to the approach embodied in the business process execution languages for Web services (WS-BPEL) [7].

The next section presents motivating scenarios and requirements for process-to-process communication. Section 3 introduces a communication model for distributed business processes addressing these requirements. Section 4 discusses the implementation and experimental evaluation of a prototype that implements these abstractions. Section 5 discusses related work and Section 6 concludes.

2 Motivating Requirements

Implementing interactions between business processes brings new issues on top of the traditional requirements found in distributed systems implementation (e.g. coupling over time and space [4,12], guaranteed delivery, encryption). In this section we distill some of these specialised requirements. These requirements are drawn from two studies on patterns in the area of integration of enterprise applications: the Enterprise Integration Patterns by Hohpe and Woolf [14], and the Correlation Patterns by Barros et. al. [9].

The correlation patterns described by Barros et. al. review conversation and message consumption patterns for distributed business processes. Most of those patterns led to requirements motivating our proposals. From Hohpe and Woolf we categorised the patterns into (1) Patterns forming requirements for a BPM messaging solution, (2) Patterns that are supported by most middleware, (3) Patterns composable from any data-aware, message-aware system, (4) Patterns concerning deployment and administration. The communication abstractions proposed in this paper focus on the first category. The technical report version of this paper [6] contains further details on this categorisation.

In the remainder of this section, we discuss the different requirements. For each requirement we cite the related patterns from [9,14].

Conversations. Support for conversations is necessary when business processes need to exchange more than one related message, and in particular when the processes are stateful, or execute over long periods. Furthermore typical process deployments, require many instances to share common channels and conversations enable messages finding their way to the correct instance. The technical report version of this paper [6] lists four forms of correlation. However this paper focuses on the most technically challenging – Property-based correlation – which assumes that many process instances share the same channel. Messages get routed to the appropriate instance by applying process-defined functions to incoming messages and then by matching the results to process instance values.

Related to: Key Based Correlation, Property Based Correlation, Reference Based Correlation, Conversation Overlap, Hierarchical Conversation, Initiator, Follower [9], Correlation Identifier [14].

Scenario 1: Purchase Order A purchase order is received. The process sends separate queries to different suppliers for each line item. Each response is correlated over the purchase order identifier, and over the line item number. This example demonstrates the need for nested conversations.

Property-Based Message Selection. Property-based message selection helps a process pick the best message off a channel. This significantly reduces the complexity logic within the process designed to iterate through an internal array of messages. There should be simpler abstractions for this.

Related to: Message Filter, Selective Consumer [14].

Scenario 2: Line Items A parts buyer wants to proceed with the best quote.

Atomic Multiple Source Consumption. When messages need to be joined, from more than one source, atomic multiple source consumption greatly reduces complexity in the process model. This is because the messages from each source may need to match over a certain field or property, and packaging them together removes the need to do this in the process. Furthermore it is much simpler to design exception handling if there is only one point of failure in the process. To demonstrate this, consider the possible outcomes of consuming them separately: (1) both message arrived, or (2) one failed to arrive, or (3) both failed to arrive, or (4) they both arrived however there was a property mismatch, or (5) there were two processes and each consumed only one message.

Related to: Atomic Consumption [9].

Scenario 3: Purchase processing Task “*ship-goods*” will not begin until there is a confirmation of credit from the accounts department, and all line items have been notified as being “in stock”.

Aggregated Consumption. Consuming messages in one “pass” of a process loop greatly reduces complexity. There is no need to iterate through a set of receive actions, or to encode loop stop-conditions, which can often be a little arbitrary if the intent is to consume all available messages. Furthermore it may be necessary to choose messages only if their properties, taken collectively, satisfy certain criteria.

Related to: Aggregator [14].

Scenario 4: Shipping Company When at least 100 items have arrived destined for the same district, and a truck is available, the truck gets dispatched.

Aggregated Consumption involving Time. Business processes are in many cases very time sensitive, for example the hours of business (9am to 5pm, Monday to Friday, EST). Thus there exists the need to include the notion of time into the process layer, for example by allowing message selection based on temporal constraints. Temporal constraints can typically be either absolute or relative i.e. “15 – 20 November, 2006”, or “within the last 7 days”. Both styles are necessary with the latter being more challenging due to the group of eligible messages being in a continual state of flux with the passage of time.

Related to: Time-Based Correlation, Moving Time Window Correlation [9]

Scenario 5: Time If, over the last five working days, more than five percent of the incoming messages that arrive at the department contain complaints then an emergency quality control process gets launched.

Contention. Contention, or competition, for the same resources, is a natural phenomenon. For instance, auctions and goods tendering rely on contention between competitors. Messages, unlike auction items, have little or no intrinsic value and therefore contention over messages may not seem compelling. However, contention over messages is an enabling technique for load balancing of message consumption; wherein many instances of the same process share the workload by only consuming messages/events when they are ready.

Related to: Competing Consumers [14].

Scenario 6: Competing Processes Copies of a process are distributed onto different hosts. The first process instance to claim the message provides the service, thus distributing the processing workload.

Interaction Cancellation. The ability to cancel incomplete interactions.

Related to: Transactional Client [14], Cancellation [2].

Scenario 7: Cancellation A supplier having posted a receive purchase order request, discovers that the warehouse needs to be replenished first, and consequently cancels its earlier receive-request.

Summary. These requirements, being drawn from the respective patterns studies, ought to be supported by most state of the art solutions, however we found that this was not the case. For instance, despite WS-BPEL being possibly the most widely accepted standard in this domain, it only provides support for the first requirement. We also found some other motivating requirements in this domain. These include event handling, channel passing, and garbage collection. These are discussed in the technical report version of this paper [6]. It is one thing to support a wide range of problems in communicating business processes, however the language constructs exposed to the creators of these processes need to be suitable to their purpose, and they should be as intuitive, or conceptually aligned to their purpose as is possible.

3 Communication Model for Business Processes

In this section we introduce a model supporting interactions between processes addressing the requirements outlined above. In the proposed model, *instances* of a process model are hosted in a *process container*. The container uses a set of channels that are referred to by the process model. The channels are hosted in some form of message-oriented middleware. Channels can be used by a process container to send messages (outbound channel), receive messages (inbound channel), or perhaps both directions (bi-directional channel).

Outbound channels are composed of a unique name, possibly a type and/or an endpoint descriptor (the latter may be determined only at runtime). The proposed model abstracts away from the specific language used to describe

message types. A particular embodiment of the concept of channel is one where message types are captured in XML Schema (or WSDL). Individual messages sent through the channel are then validated against its XML schema (or WSDL message type definition). Similarly, we abstract away from the mechanism used to describe destination endpoints. If using HTTP as a transport protocol for example, an endpoint can be described as a URL, while if using SOAP/HTTP, it can be described as a WSDL operation binding. Outbound channels support a range of message sending primitives which are described in detail in [5]. In the rest of the paper however, we focus on inbound channels.

Inbound channels have the same components as outbound ones but they additionally have a set of *properties*. A property is a function that takes as input a message and produces a literal value. This is similar to the concept of *property alias* defined in BPEL. However, as shown later, the scope of applicability of properties in our model is wider than that of BPEL. In BPEL, property aliases (and their composition in the form of *correlation sets*) are only used to correlate pairs of outbound and inbound messages. In contrast, in our model, properties can be used to perform other forms of message selection and aggregation.

A relation (i.e. a database table) is created for each inbound channel used in any process model. Each relation contains two predefined attributes: one of type **message identifier** and the other of type **timestamp**. Additionally, the relation contains one attribute (column) per property associated to the channel.

Properties are used to define *filters*. A filter is a function that is evaluated against the set of messages available for consumption over one or multiple channels. When the evaluation of a filter returns a non-empty set of messages, we say that the filter *matches* these messages. Filters can fulfill two purposes: (i) *garbage-collecting filters* are registered by a business process to discard unwanted or unnecessary messages over inbound channels; (ii) *message consumption filters* are used to consume one or multiple messages from one or multiple channels. Orthogonally, filters may be *one-off* or *persistent*. A one-off filter is immediately withdrawn after it has matched a message or set of messages, while persistent filters are preserved until explicitly withdrawn.

A filter is represented as a query over the relations(s) associated with the channel(s) it refers to. These queries are always constrained to produce a relation wherein each tuple contains **message identifier** attributes.

When a message arrives onto a channel, a tuple is inserted into the relation associated with that channel. This tuple always contains **message identifier** and **timestamp** attribute values, as well as attribute values obtained by applying each of the channel's properties to the incoming message. The inserted tuple represents the incoming message for the purpose of evaluating message filters. After being abstracted as a tuple, the incoming message is either:

- Immediately routed to a message receipt action if one has registered a filter that matches the message (possibly in combination with other messages).
- Discarded if the message matches any of the garbage-collecting filters registered for that channel.
- Queued until it matches a garbage-collecting or message consumption filter.

Conceptually, a filter is re-evaluated every time that a new message arrives to any of the channels it refers to (or continuously in the case of filters whose query depends on the current time). In practice however, the evaluation can be made incrementally and only when required.

Primitives for registering and withdrawing filters are provided as part of the communication framework. Registration and withdrawal of filters can be initiated either by the process container or by individual process instances. When a message consumption filter is registered, the filter is run once against the set of messages available in the channels referenced by the filter. If the filter matches one or several messages, these are removed atomically from their channel(s) and given back to the process container or process instance that registered the filter. If the filter is one-off, it is withdrawn. Should no match between a filter and the existing set of messages be found upon registration of the filter, the filter is maintained and re-evaluated whenever required as explained above, until the filter is either explicitly withdrawn or it matches a message (or set of messages). Once a match is found, the message(s) are routed to the corresponding process container or instance and the filter is removed.

Garbage-collecting filters work similarly: when registered, the filter is evaluated against the contents of the channels targeted by the filter. If a match is found, the matched message(s) are discarded. If the filter is one-off, it is withdrawn otherwise, it is preserved and it is re-evaluated when required.

Fig. 1. Petri net capturing the treatment of inbound messages

The proposal is formally captured as a Coloured Petri net in Fig. 1. The net shows how inbound messages are stored and matched against filters. A token in place “*incoming message*” represents a message received by the communication layer. A transition called “*put*” moves tokens from this place to a place called “*message buffer*”. This latter place holds a single token containing a list of all unmatched messages over all channels. Two of the places “*garbage collecting filters*” and “*message consumption filters*” are meant to contain one token per active filter. Transition “*collect garbage*” fires when there is a garbage-collecting filter that matches at least one of the messages in the message buffer. This transition puts back a modified message buffer in which all messages matching

the garbage-collecting filter have been removed. Similarly, transition “*match*” fires if there is a consumption filter matching at least one message in the buffer. This transition also puts back a modified message buffer in which the matched messages are removed. In addition, it produces a tuple containing the filter and the set of matched messages into output place “*matched messages*”. These tokens can then be routed to the process container or process instance that registered the filter in question. The latter is identified by a *request identifier* (“*reqID*”). For simplicity, the net only captures the case of “persistent filters”, but it is easy to extend the net to deal with one-off filters: the only difference being that such filters should not be put back by transitions “*collect garbage*” and “*match*”.

The proposed communication model abstracts away from the way channels and filters relate to business process activities or events. This way, the model can be integrated into a wide range of process definition languages. In BPEL, for example, inbound communication actions appear in two forms: as a standalone *receive* activity type and as the second leg of activities of type *invoke*, where the first leg corresponds to an outbound communication action. Thus, BPEL can be extended with the proposed communication primitives by enabling *receive* and *invoke* activities to refer to channels and filters as defined above. Channels can then be linked to *partner links* and *operations* in BPEL.

Similarly, the proposed model can be used to extend the YAWL process definition language to support a richer set of communication patterns. For example, we can define a type of message receipt task in YAWL such that: (i) upon enablement, the task registers a one-off message consumption filter defined as part of a task decomposition; (ii) the task then waits until the filter returns a match; (iii) should the task be cancelled before a match is found, the filter is withdrawn. Also, we can allow message consumption filters to be attached to the *initial condition* of a YAWL process model, to capture scenarios such as: “a new process instance should be started whenever a given type of message (or combination of messages) has been received.” An integration of the proposed communication primitives into YAWL is left as future work.

4 Implementation and Evaluation

The implementation builds on a middleware service, and API, called JCoupling (available from www.sourceforge.net/projects/jcoupling). JCoupling supports a superset of the communication styles supported by mainstream communication middleware. It also abstracts away from transport protocol details, thus allowing us to concentrate on the core aspects of our proposal. It supports uni- and bi-directional communication, time, space, and synchronisation decoupling, and provides support for fault propagation. JCoupling can operate over open-JMS² and is able to use XML/HTTP or XML/TCP for message transport.

Figure 2 presents an architectural diagram of the prototype. It shows properties and filters being used during a simple interaction between two process tasks. In Fig. 2: (1) A message is sent by task “*T1*” over channel “*Ch1*” (denoted by

² Open JMS www.openjms.sourceforge.net accessed November 2006.

arrows labelled “1” going from the engine to the controller). (2) Properties “ $P1$ ” and “ $P2$ ” are used to extract values from the message. Next the message is put onto JCoupling (3.a), and a new tuple (row) is added to the relation (table) for channel “ $Ch1$ ” (3.b). (4) The receiver task “ $T2$ ” posts a filter over channel “ $Ch1$ ”. (5) Using the filter, the controller performs a query, over the relation for channel “ $Ch1$ ”. That query produces a tuple and the matching message is extracted from JCoupling (6). The callback to “ $T2$ ” contains the message (7).

Fig. 2. Architecture of the proposal

This section gives further details on the implementation of properties, interactions and filters, and reports on an experimental evaluation of the prototype.

4.1 Implementing Properties

The prototype implementation contains an interface called **Property**. This interface and three implementing classes are presented in Fig. 3. A complete listing of this interface, and an explanation of its concrete classes is provided in the technical report [6]. Instances of **Property** extract various scalar values from messages.

We envision a configuration tool for creating property instances such that process designers do not need to write Java code. For example if the process designer wishes to create an XPath property, to extract “*PurchaseOrderID*’s” from messages then it would only be necessary to supply a property name, a channel binding, an XPath expression, and a Type (e.g. ‘PurchaseOrderId’, ‘PO_Chann’, ‘/order/@po-id/text()’, ‘Text’).

As each new message arrives the prototype adds a tuple to the *property relation* for that message’s channel using a PostgreSQL³ database. As mentioned

³ The PostgreSQL database system www.postgresql.org (accessed February 2007).

Fig. 3. UML of the property Interface

in Sect. 3, each channel’s property relation has two default attributes (columns): `messageid` and `timestamp`. Table 1 presents a relation for channel `QuotesCh`.

Table 1. Relation corresponding to channel `QuotesCh`

QuotesCh			
messageid	timestamp	quote	amount
C0000M0001	2006-11-15 16:35:50	4000	1100
C1111M0002	2006-11-16 18:12:20	3000	1000
C2222M0001	2006-11-16 20:42:53	2500	1001
C3333M0003	2006-11-17 16:57:21	2000	999

4.2 Resolving Contention

A process may have to wait days before an appropriate message can be received. Hence, the prototype stores send/receive requests and performs callbacks when the interactions are complete. Indeed requests to send/receive, have their own lifecycle, including interaction cancelation (i.e. when process state change makes unnecessary an incomplete interaction). A race condition between two receive tasks is a perfect example of this. For example, a task “receive-bill-payment” and a task “receive-purchase-order-cancelation” operate such that the completion of one disables the other. The first message to arrive will make the other message unnecessary. This is a “pick” activity in BPEL, or a “deferred choice” in YAWL [1]. Accordingly the prototype (being a control layer between a process and the message layer), exposes a ‘withdraw’/‘cancel’, primitive for send and receive interactions – also implementing the cancelation requirement (Sect. 2).

Another form of contention occurs where two tasks both want the same message/s. Contention may be a requirement (for example load balancing, Scenario 6), or it may be accidental, or perhaps even unavoidable due to the nature of the business process. Regardless contention, being a natural and sometimes necessary phenomenon, mandates a graceful approach to handling it⁴.

⁴ This is a distinguishing point from WS-BPEL, which throws runtime exceptions when contention between two receive-tasks occurs. Also, WS-BPEL’s greedy routing of messages to process instances precludes the possibility to support contention.

To overcome the problems of contention within the context of our proposal we adopted the following algorithm in the prototype.

1. The receive request, containing its messages filter, is stored in the prototype.
2. Once the filter produces a non empty set (Ω) of message-identifiers, it is locked to ensure that the receiver (task) cannot withdraw the request.
3. Then the message engine attempts to lock every message in Ω .
4. If all messages were successfully locked, then for each message-id in Ω , that property tuple is deleted, and each corresponding message gets removed from the buffer and sent to the requestor in a callback.
5. Finally the request filter is withdrawn from the engine.

Should the receiver be unable to lock the filter at step 2 the prototype withdraws the filter. At step 3, should any of the messages in Ω be already locked then the filter locked during step 2 is unlocked, and the filter is rescheduled. Hence the first locker of the message succeeds.

4.3 Implementing Filters

This section shows how the extensions to JCCoupling outlined above, support the motivating requirements of Sect. 2. Property relations, abstracting from message content, enable the use of restricted SQL queries to produce relations containing message identifiers. The primary restriction we place over the queries is that their outer projections must only be over attributes of the domain `messageid`. i.e. for all attributes of a relation produced by applying any query to the property relations, the domain of that attribute must be `messageid`.

A single property, or combinations of properties can be used to select messages. We envision that the process modeller would have a library of configurable *filter templates*, each containing semi-complete queries. The filter templates would also allow a “design time” binding to process variables, enabling runtime data to be inserted into parameterised SQL. Alternatively, for those situations requiring sophisticated aggregate operations, or complex joining expressions, the process designer may instead write their own SQL.

Property Based Selection. Scenario 2 captured the need to select and proceed with the best quote, which is an example of selecting messages based on property values. The query in Listing 1 uses two properties defined over the channel “QuotesCh” (see Table 1). These are “price” and “quantity”. When a receive request containing a query is invoked over the prototype, it will apply the query – returning messages in a callback to the requestor when results are found.

Listing. 1. This query combines ‘price’ and ‘quantity’ to find the best value offer

```

1 SELECT messageid
2 FROM QuotesCh
3 WHERE quantity >= 1000
4 AND price/quantity =
5 ( SELECT min (CostPerUnit)
6   FROM ( SELECT price/amount AS CostPerUnit
7         FROM QuotesCh ) AS Q1 )

```

Conversations. Scenario 1 outlined the need to correlate messages to an outer conversation for “purchase-order-ID” and an inner (nested) conversation for “line-item-id”. Messages will only be correlated to a nested conversation if they satisfy correlation filters for the inner conversation, and all parent conversations. To achieve this we append an **AND-Clause** to the query. The following query (see Scenario 1) extracts runtime data from two process variables visible to the receive task: namely `PurchaseOrderID` and `LineItemID`.

Listing. 2. Achieving nested correlation through querying correlation properties

```

1 SELECT messageid
2 FROM PoResponseCh
3 WHERE PoID = $PurchaseOrderID$
4 AND ItemID = $LineItemID$

```

We envision that the process layer will generate queries for conversations from a “conversation” construct in the process model. This construct would declare which message properties are to be used for correlation, and whether each communication task involved initialises the conversation, or follows it. Furthermore if any task involved in the conversation wants to apply message filters we may have the task append more **AND** clauses to the generated query.

Atomic Multiple Source, and Aggregate Consumption. Achieving a combination of atomic multiple source consumption and aggregate consumption is possible by applying a query to two or more property relations. For example Scenario 4 required aggregate consumption of 100 packages destined for the same area, and a truck availability message from another channel. Listing 3 extracts runtime data from a process level variable called `deliveryDistrict`. The query either returns at least 100 tuples, or returns nothing.

Listing. 3. This query produces a relation of `messageid`’s wherein the attribute (`Pack.messageid`) refers to messages from Channel `Packages`, and the attribute `Truck.messageid` refers to a message on Channel `TruckWaiting`. Related messages are linked, thus removing the need to relate messages in the process.

```

1 SELECT Pack.messageid, Truck.messageid
2 FROM Packages As Pack, (SELECT messageid FROM TruckWaiting LIMIT 1) AS Truck
3 WHERE Pack.deliveryDistrict = '$deliveryDistrict$'
4 AND 100 <= ( SELECT count(*)
5              FROM Packages
6              WHERE deliveryDistrict = '$deliveryDistrict$' )

```

Aggregated Selection Involving Time. Scenario 5 sought to find if at least 5% of messages that are less than 7 days old contain complaints; drawing on solutions to both time and aggregated selection. Additionally complaints were sought on *all channels*. A join between property relations will not work as every tuple from each relation represents one discrete event that needs to be considered separately. So unlike Listing 3, the result relation will have one attribute. A **view** over the **union** of source channels (property relations), solves this. In Listing 4 this view is named “Merged”. It combines properties (`messageid`, `timestamp`, and `complaint`) from each of the source property relations.

Listing. 4. Using `union` and `view` to combine input sources for aggregate operations

```

1 CREATE VIEW Merged AS (
2   SELECT messageid, timestamp, complaint
3   FROM CustomerCh
4   UNION SELECT messageid, timestamp, complaint
5   FROM PartnerCh )

```

Using the above view, aggregate calculations over the messages, taken collectively, becomes feasible. Listing 5 demonstrates this. Lines 5 – 12 produce `false` unless 5%, or more, of the messages are complaints. Lines 3, 8, & 12 show the use of relative time expressions over the `timestamp` property. In cases like this we imagine that the process creator would write Listings 4 and 5 manually.

Listing. 5. This query, adapted from a continuous query in [8], produces a non-empty result of `messageid`'s when 5%, or more, of last week's messages contain complaints

```

1 SELECT messageid
2 FROM Merged
3 WHERE timestamp > (CURRENT_TIMESTAMP - INTERVAL '7 days')
4 AND complaint = true
5 AND ( SELECT count(*)
6       FROM Merged
7       WHERE complaint = true
8       AND timestamp > (CURRENT_TIMESTAMP - INTERVAL '7 days')
9     ) >= (
10      SELECT 0.05 * count(*)
11      FROM Merged
12      WHERE timestamp > (CURRENT_TIMESTAMP - INTERVAL '7 days')
13    )

```

4.4 Performance Evaluation

To compare our correlation approach with that of WS-BPEL, we conducted experiments where up to ten thousand interactions were executed over our prototype and over a WS-BPEL simulator. Each experiment involved creating, in random order, at fixed intervals, a set of XML messages, all identical except for the value of one element which was mapped to a property. Receiving processes were spawned in the same way, each of which waited for one of the created messages. The code for the experiments is released with the JCCoupling distribution.

The WS-BPEL correlation simulator was built using the same middleware as our prototype. This simulator receives messages off a designated channel. Each time a message arrives, an XPath expression is evaluated against it to extract a property value (this is called a `propertyAlias` in WS-BPEL). The extracted property value is then stored in a hash table together with the corresponding message identifier.⁵ Concurrently, the simulator handles requests to consume incoming messages based on property values. When the simulator finds a match between a receive request and a message, the corresponding entry is deleted from the hash table, symbolising that the message has been correlated.

The test results are presented in Table 2. The table shows that our approach slightly underperforms that of WS-BPEL for small numbers of messages. This

⁵ In the interest of fairness, each entry is written to disk after being added to the in-memory hash table, as our prototype stores property values in persistent tables.

Table 2. Results of performance tests

Number of Interactions		50	100	250	500	1000	2500	5000	10000
Time (ms)	Proposed Approach	611	951	2230	5064	9003	25276	71409	234211
	WS-BPEL Approach	524	938	2103	4046	7825	20506	44204	133933
Performance Difference		14%	1%	6%	20%	13%	19%	38%	43%

difference grows for larger numbers of messages. The accentuated difference can be explained by the fact that in the implementation of our approach, queries to match uncorrelated messages with pending receive requests are run against a persistent database, whereas in the WS-BPEL simulator, the corresponding lookup is done in-memory⁶. For larger numbers of messages, this leads to a performance penalty due to database cache management. A more consistent performance could be achieved by using an in-memory database system to implement our approach. Indeed, it is not necessary to make the correlation data structures persistent, only the messages themselves need to be persistent. In future, we plan to implement a more refined version of our approach and run a fairer and more detailed comparison against the WS-BPEL approach.

The performance penalty of our approach should be weighed against the additional functionality that it brings in. Indeed, our approach supports aggregate messaging, multi-source consumption and message contention. Moreover, as previously mentioned, there are opportunities to optimise the brute-force approach used in our implementation through incremental query evaluation.

5 Related Work

Communication in the context of distributed business processes has traditionally been researched from the perspective of *protocol* or *contract* definition. For example, the CrossFlow system [13] enables process designers to define contracts governing the communication between multiple workflows, possibly distributed across organisational boundaries. These contracts can be statically checked for consistency. Similarly, [3] proposes a method for capturing inter-workflow communication protocols and detecting deadlocks that can arise when inter-connecting processes with incompatible communication protocols. This body of work is complementary to our proposal, as we do not deal with static analysis, but rather with the routing of messages to processes at runtime.

WS-BPEL exhibits strong support for conversations with the exception of *Instance Channels*, which are not supported because WSDL *ports* are not generated during process execution. WS-BPEL, does not support the selection of messages based on their properties despite the use of properties in correlation sets. Nor does it support atomic-batched consumption. Contention over messages is not supported as messages are greedily consumed off channels and allocated to process instances immediately. Time-based message selection is not supported either. A WS-BPEL process consumes everything sent to it, thus forcing the modeller to select and throw away unwanted messages, as part of the core process

⁶ Although entries are written to disk, lookups over the hash table are done in-memory.

logic. WS-CDL [15] has some distinguishing features, with respect to WS-BPEL, such as abstractions for channel passing and a global viewpoint over all actors in a choreography. However, in terms of the motivating requirements outlined in this paper, WS-CDL has very similar strengths and weaknesses to WS-BPEL.

Widom et al. [8,16] propose an approach to optimising the evaluation of continuous queries over one or many data streams. They address some of the problems associated with scalability of such queries and propose incremental evaluation techniques based on the type of query. We plan to apply some of their findings to enhance and optimise the evaluation of filters used in our proposal.

6 Conclusion

This paper proposed an inter-workflow communication and control layer, to lie between traditional workflow and messaging layers – providing an isolated area for the description and execution of communication. This would provide relief from using “spaghetti” solutions to achieve non-trivial interactions. The proposal is based on a strong and relatively simple set of abstractions – namely channels, properties, property relations and filters. Channels abstract from middleware topics, and queues. Properties abstract from message content and format, while property relations provide the foundation for property filters. Filters abstract from the business level requirements for choosing and selecting messages. These enable all forms of correlation, message selection, aggregated message consumption, and time based message consumption, over a single or multiple channels evaluated collectively. The possibility of contention between process/task instances is overcome by locking messages, and filter requests. The proposal has been implemented on top of a communication API, namely JDecouple, that abstracts away from the underlying middleware and communication protocols.

Future work will aim at integrating the proposed communication abstractions into process definition languages. Specifically, we aim to extend YAWL with the ability to attach communication actions to various elements of the notation.

Acknowledgements. This work is funded by ARC Discovery Grant DP0451092. The third author is funded by a Queensland Smart State Fellowship.

References

1. van der Aalst, W., ter Hofstede, A.: YAWL: Yet Another Workflow Language. *Information Systems* 30(4), 245–275 (2005)
2. van der Aalst, W., ter Hofstede, A., Kiepuszewski, B., Barros, A.: Workflow patterns. *Distributed and Parallel Databases* 14(3), 5–51 (2003)
3. van der Aalst, W., Weske, M.: The P2P approach to Interorganizational Workflows. In: Dittrich, K.R., Geppert, A., Norrie, M.C. (eds.) *CAiSE 2001*. LNCS, vol. 2068, pp. 140–156. Springer, Heidelberg (2001)

4. Aldred, L., van der Aalst, W., Dumas, M., ter Hofstede, A.: On the Notion of Coupling in Communication Middleware. In: In Proceedings of the 7th International Symposium on Distributed Objects and Applications (DOA). Agia Napa, Cyprus, November 2005, pp. 1015–1033. Springer, Heidelberg (2005)
5. Aldred, L., van der Aalst, W., Dumas, M., ter Hofstede, A.: Understanding the challenges in getting together: The semantics of decoupling in middleware. Technical Report BPM-06-19, Business Process Management Center, Brisbane, Qld, Australia, 2006 ((February 2007), <http://www.bpmcenter.org> accessed
6. Aldred, L., van der Aalst, W.M.P., Dumas, M., ter Hofstede, A.H.M.: Abstractions for communication between distributed business processes. Technical Report BPM-06-28, Business Process Management Center, Brisbane, Qld, Australia, 2006, accessed (February 2007) www.bpmcenter.org
7. Alves, A., Arkin, A., Askary, S., Bloch, B., Curbera, F., Goland, Y., Kartha, N., Liu, C., König, D., Marin, M., Mehta, V., Thatte, S., van der Rijn, D., Yendluri, P., Yiu, A.: Web Services Business Process Execution Language. Initial Draft of standards proposal by OASIS, June 2006, accessed (July 2006) <http://www.oasis-open.org/apps/org/workgroup/wsbpel/>
8. Babu, S., Widom, J.: Continuous queries over data streams. SIGMOD Record 30(3), 109–120 (2001)
9. Barros, A., Decker, G., Dumas, M., Weber, F.: Correlation Patterns in Service-Oriented Architectures. In (FASE). Proceedings of the 10th International Conference on Fundamental Approach to software Engineering, Braga, Portugal, March 2007, pp. 245–259. Springer, Heidelberg (2007)
10. Beugnard, A., Fiege, L., Filman, R., Jul, E., Sadou, S.: Communication Abstractions for Distributed Systems. In: Buschmann, F., Buchmann, A.P., Cilia, M.A. (eds.) Object-Oriented Technology. ECOOP 2003 Workshop Reader. LNCS, vol. 3013, pp. 17–29. Springer, Heidelberg (2004)
11. Cypher, R., Leu, E.: The semantics of blocking and nonblocking send and receive primitives. In: Siegel, H (ed.) Proceedings of 8th International parallel processing symposium (IPPS), pp. 729–735 (April 1994)
12. Eugster, P., Felber, P., Guerraoui, R., Kermarrec, A.: The Many Faces of Publish/Subscribe. ACM Computing Surveys 35(2), 114–131 (2003)
13. Grefen, P., Aberer, K., Hoffner, Y., Ludwig, H.: CrossFlow: Cross-organizational Workflow Management in Dynamic Virtual Enterprises. International Journal of Computer Systems, Science, and Engineering 15(5), 277–290 (2001)
14. Hohpe, G., Woolf, B.: Enterprise Integration Patterns: Designing, Building, and Deploying Messaging Solutions. Addison-Wesley, Boston (2003)
15. Kavantzaz, N., Burdett, D., Ritzinger, G., Fletcher, T., Lafon, Y., Barreto, C.: Web Services Choreography Description Language Version 1.0. Candidate Recommendation, (November 2005) <http://www.w3.org/TR/ws-cd1-10/>
16. Olston, C., Jiang, J., Widom, J.: Adaptive filters for continuous queries over distributed data streams. In: SIGMOD '03: Proceedings of the 2003 ACM SIGMOD international conference on Management of data, pp. 563–574. ACM Press, New York (2003)