

Consistency between Executable and Abstract Processes

Axel Martens
IBM T. J. Watson Research Center,
Component Systems Group
E-mail: amarten@us.ibm.com

Abstract

Process models play an all-important role in the development of cross-organizational business processes. On the one hand, the interaction between the participating companies often is specified globally, for example by means of multiple abstract process models – one for each partner. On the other hand, each partner defines its local process autonomously in terms of an executable process model. The important question is whether such an executable model is consistent to the predefined abstract model. This paper describes an approach to prove this property automatically.

Emerging technologies and industrial standards in the field of *Web services* enable a much faster and easier cooperation of distributed partners. This paper is concerned with the application of Web services to distributed, cross-organizational business processes.

A *Web service* [1] is a self-describing, self-contained modular application that can be described, published, located, and invoked over a network, e.g. the Internet. A Web service performs an encapsulated function and can be accessed via a standardized interface. In this paper, each local sub-process of each participating company is realized through one Web service. The composition of all Web services of all participating companies realizes the global business process.

The Web service approach provides a stack of closely related technologies [3] to cover heterogeneity and distribution underneath a homogenous concept of components and composition. Among other things, the *Business Process Execution Language for Web Services* BPEL4WS [2] belongs to this stack. Due to this layered architecture, the presented method can be focussed on the Web service's process model in terms of BPEL4WS – without losing generality or practical relevance.

The goal

Process modeling is one of the most crucial tasks while developing or integrating enterprise applications. In practice, there are various levels of abstraction to

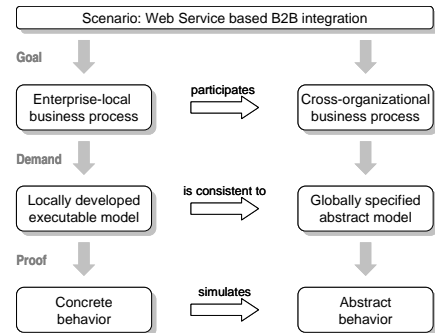


Figure 1. Field of application

model a business process: A coarse grained model visualizes the overall dependencies between the activities. Usually, it is refined in several steps during the development process. Finally, a very precise model, covering all necessary aspects, can be executed directly by help of a workflow engine. Hence, such a model is called *executable* model, whereas each model of a prior stage can be called “*abstract*” model in some sense. BPEL4WS can be used to express both: the abstract process model as well as the executable process model.

Referring to Web service based business processes, for each local sub-process, often there is already one executable BPEL process model. But, the distributed global process itself has to be modeled, as well. Such a *global model* is required, especially, to ensure the compatibility or to specify the interaction in terms of a contract between the partners. And, although the global process is realized through composition of multiple local sub-processes, to establish the global model, however, it is not adequate to simply combine the local *executable* models of the participating partners (e.g. to keep business secrets). Instead, the global model should aggregate one *abstract* model for each partner.

If both models are given a certain partner, an important question is whether the executable model is *consistent* to the abstract model. Consider for example Web portal, which is offering various online shopping

services to potential customers. This portal has defined an abstract process model such that each customer can structure his own process accordingly to this model. If a new shop wants to join this portal, its process should interact with each customer properly, i.e. the shop's executable process model *has to be* consistent to the predefined abstract process model – Figure 1 shows this correlation. This paper describes an approach to verify the consistency between process models automatically.

The method

As this paper will show, very differently structured process models might be called consistent. Hence, the presented method focusses on the *behavior* rather than the structure: An executable model is called to be *consistent* to an abstract model if and only if the concrete behavior *simulates* the abstract behavior (cf. Figure 1). By this means, each component that was built according to the abstract model can't see any difference while interacting with the executable process.

The analysis of consistency runs in three steps. First, a precisely and explicitly defined semantics of the modeling language is needed. This approach is based on *Petri nets* [17], because this formal method is widely used for modeling and analyzing business processes and Web services [24, 4, 5]. The second step extracts the relevant information and generates the *communication graph* [11] – an explicit representation of the process' external visible behavior. Rest upon this necessary abstraction, the third step verifies the *simulation* between concrete and abstract behavior by comparing the corresponding communication graphs. All three steps are performed automatically by the tool WOMBAT4WS [29]. Hence, the user has to deal with BPEL process models, only. Moreover, the method could easily be extended to other modeling languages (e.g. WS-CDL [6], OWL-S [27] or YAWL [22]), by adjusting the first step only.

The paper is structured according to Figure 1: After the goal was described, Section 1 delves into the already mentioned use case of the Web portal, and presents *abstract* and the *executable* BPEL process models. Section 2 sketches the extraction of the *behavior* of a BPEL process models and discusses different ways to compare behavior. The proposed notion of *simulation* is presented in Section 3. Finally, Section 4 summarizes the results, points to additional use cases, and relates the work to other published approaches.

1 Examples

A Web portal offering customers access to various online is used as running example through this paper. The language BPEL4WS allows to define abstract process models as well as executable models. Basically, an

executable model is required to be completely specified, e.g. all variables have to be declared and used properly. In contrast, an abstract model may leave out some of those aspects (cf. [8]). Although this paper refers to these notions, it focusses on the semantic distinction rather than the syntactic difference: An executable models contains more details than the abstract model does. Hence, the notion 'executable model' could be replaced by the notion 'refined model', if preferred.

1.1 The abstract process model

The business strategy of the Web portal has two aspects: On one hand, the portal is open to all online shops. On the other hand, the portal requires the participating shops to build their services according to a standardized protocol, specified in terms of an abstract BPEL process model. Figure 2 shows this model.

Basically, a BPEL model consists of two different kinds of activities: *Basic activities* are used to communicate to the outside (e.g. receive order, send questions in Figure 2), to perform internal steps (e.g. process payment) or to interfere with the control flow (e.g. by signaling fault and process termination – not present in the example). *Structured activities* aggregate other activities. Therefore, they are used to build the control structures of the process, i.e. alternative branches – either based on incoming messages (*customers payment*) or based on data (*chose shipment*) – loops (*while order incomplete*), sequential concatenation (*process activity*), and parallel branches (*process activity of Figure 3*).

The process shown in Figure 2 models the following behavior: First, the customer should place his order. Then, the shop may send zero or more questions to the customer concerning his order and await his update. Eventually, the shop sends the invoice and requires the customer to pay, either with credit card (*cc payment*) or out of his checking account (*ca payment*). Finally, the shop finishes the process by sending the *delivery data* or *pickup data* according to its business strategy. The formats and the channels of messages being exchanged are defined in the *Web service description WSDL* [1].

1.2 The executable process model

Figure 3 shows the executable process model of an online shop that wants to participate in the Web portal. Compared to the abstract model, this process follows a different strategy. First, it offers to the customers a product query that may precede the actual order, and second, no pickup of ordered products is available. Due to this business strategy, the structure is quite different compared with the abstract process. Nevertheless, this process model can be proven to be consistent.

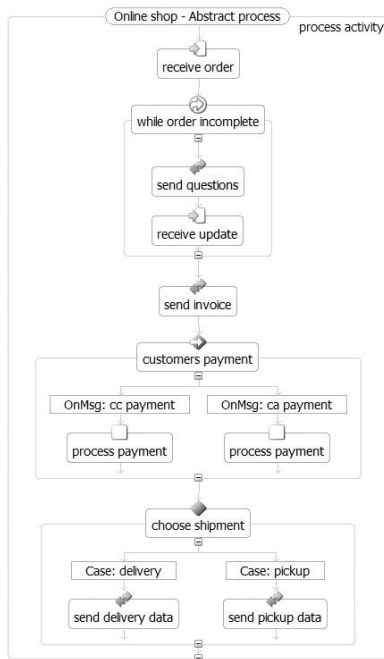


Figure 2. Abstract BPEL process model

The executable process is structured into two concurrent activities: *customers initial choice* and *order processing*. But, these activities are synchronized by two *links*. Hence, the *order processing* is started only if the customer send an *order* right at the beginning or he sends a *request*, the product is available, and he decides to order after he got the *confirmation*.

Assuming the customer has ordered a product, he either gets the invoice or he is asked *questions* concerning his *order* exactly once. In contrast, the abstract process model allows an arbitrary number of callbacks. The payment is handled the same way in both processes. Finally, the executable process sends the *delivery data*.

In spite of the many differences, it is not difficult to see: Each customer, who behaves properly according to the abstract process (i. e. behaves like an *utilizing environment*, cf. Section 2.1), has no problems to cooperate with the executable process. But there is no obvious way to prove this property just by syntactical matching. In fact, the executable process follows a common business strategy for smaller companies: On the one hand, it respects the rules made by a company with larger market power, but on the other hand, it offers additional services to attract more customers. Hence, it is worth supporting to this use case methodically.

2 Behavior

The *Business Process Execution Language for Web Services* BPEL4WS [2] is in the very act of becoming the industrial standard for modeling Web service based

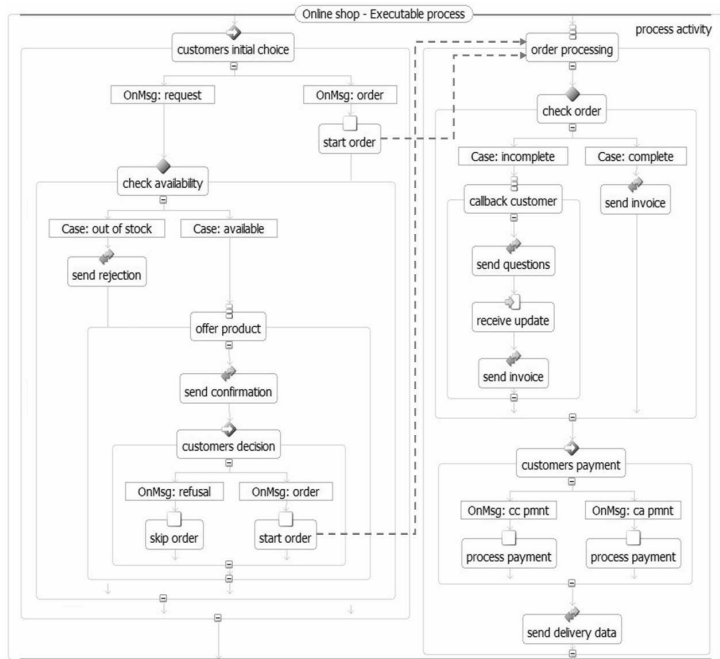


Figure 3. Executable BPEL process model

business processes. But, its semantics so far is defined only by English prose or encoded into middleware components, more ore less accurately. To verify properties like properties like the consistency of two models a formal semantics of all its concepts is needed. The presented method is based on a Petri net semantics [20].

2.1 Modeling with Petri nets

Petri nets are a well established method for modeling and analyzing (cross-organizational) business processes [23, 25, 7]. Moreover, beside the semantics for BPEL4WS, other recent research projects apply Petri nets to Web Services [5, 16]. Because of the limited space, only a very brief overview is given here.

Figure 4 shows four small BPEL processes and their corresponding Petri net models. A *Petri net* $N = (P, T, F)$ consists of a set of *transitions* T (boxes), a set of *places* P (ellipses), and a *flow relation* F (arcs) [17]. A transition represents a dynamic element, i. e. an activity of a business process. A place represents a static element, i. e. the causality between activities or a message channel. The state of a Petri net is represented by black *tokens* distributed over the places.

A Web service defines an internal structure and an interface to communicate with other Web services. Hence, the Petri net model of such a Web service consists of a *workflow net* [24] supplemented by set of places representing directed message channels. Such a model is called *workflow module* [13]. Two modules are composed by merging all common interface places

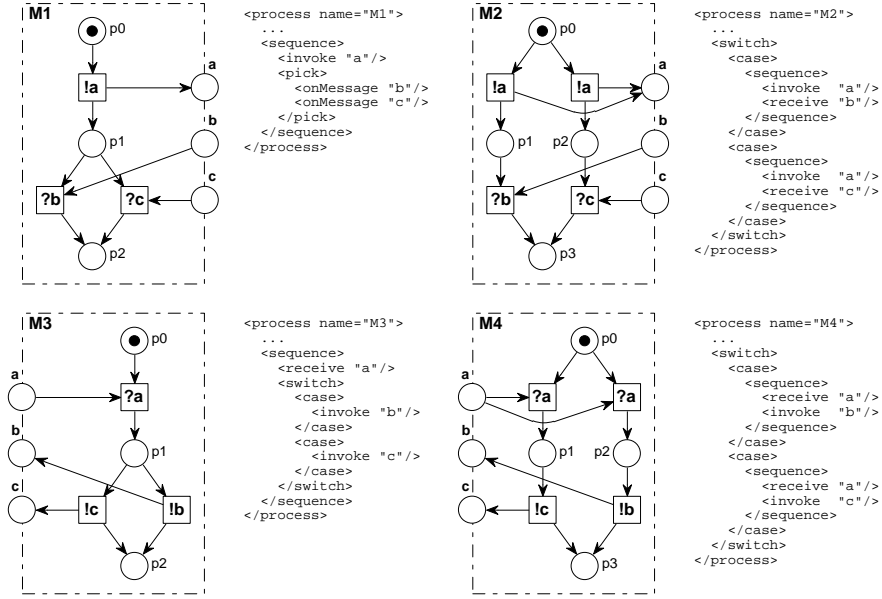


Figure 4. Examples of (non-)equivalent modules

(matching by name). The dangling input and output places become the new interface of the composed module. Additionally, some net elements for initialization and termination are added. If both sets of interface places completely match, one module is called *environment* of the other (obviously, this notion is symmetric, cf. Figure 4). A Module M is called *usable* if there exists at least one environment U such that the composed system $M \oplus U$ is a sound process model (i. e. the Petri net should be *weak-sound* [13]). Such an environment is called an *utilizing environment* (cf. Section 2.2).

Due to the mapping into Petri nets, several analysis methods are now applicable to BPEL processes models: the verification of *usability* of one Web service [11], the verification of *compatibility* of two Web services [10], the *automatic generation* of an abstract process model for a given Web service [12], and the verification of *simulation* and *consistency* – presented in this paper.

2.2 Comparing behavior

Based on the underlying Petri net formalism, it is now possible to reason on the behavior of a BPEL process with well defined notions of current state, enabled activities and occurrence of transitions. Intuitively, one process simulates another process if in a *comparable situation* the first process *behaves like* the second process. Concerned with various formal methods, there are countless approaches published dealing with the comparison of behavior in terms of simulation or equivalence. This variety results from different interpretations of the terms “comparable situation” and “behave like” ([26] gives a substantial overview on equivalence

notions). But, none of them seems to fit exactly to this field of application, i. e. none of them allows an executable process to accept more messages and send less messages (compared to the abstract process). To give an idea, the following paragraphs compares two of them, and presents the preferred definition at the end.

Trace equivalence: Two communicating systems could be compared by looking from the outside only at the sequences of messages being exchanged – called *traces*. From this point of view, the modules M1 and M2 both simulate each other, because they have identical traces a,b and a,c. But, module M3 is an utilizing environment of module M1, whereas the composition of M3 and module M2 yields a deadlocking system. Moreover, it is possible to prove the non existence of any utilizing environment for module M2. Hence, the requirement of trace equivalence is too weak for our purpose.

Bisimulation equivalence: The opposite approach compares the internal states and enabled transitions of both modules: A state z_s simulates a state z if each transition (identified by its label), which is enabled in z , also is enabled in z_s , and the reached state z'_s simulates the reached state z' . Two modules are equivalent if their initial states simulate each other. In that sense, module M4 does not simulate module M3, because neither state [p1] nor state [p2] of module M4 simulates the state [p1] of M3. But, module M1 is an utilizing environment of module M3 and an utilizing environment of M4 as well. Moreover, it is possible to prove that each environment utilizes either both modules or none. Hence, the requirement of bisimulation equivalence is a too restrictive for our purpose.

This short discussion shows that an adequate notion of simulation (or equivalence, resp.), first of all, should be derived semantically from the field of application. In a second step, a proving method can be developed.

Definition 2.1 (Simulation/Equivalence).

A workflow module A *simulates* a workflow module B if each utilizing environment of module B is an utilizing environment of module A , too. Two workflow modules A and B are called *equivalent*, if module A simulates module B and module B simulates module A . *

This definition exactly meets the requirement of the cross-organizational business process scenario: Lets consider a workflow module M and an utilizing environment E . For each module M' simulating M and for each module E' simulating E holds: E' is an utilizing environment of M' and vice versa. This property follows directly from Definition 2.1. The remainder of this paper describes the verification of simulation.

2.3 Reflecting behavior

As discussed in the previous paragraph, a level of abstraction is needed that still combines the internal and the external views on a workflow module. A Web service is a reactive system, it consumes messages from the environment and produces answers depending on its internal state. But, an environment, which wants to use this service, has no *explicit* information on the service's internal state. Instead, each environment can derive some *implicit* information by considering the communication towards the module. Exactly that kind of information is reflected within a data structure – called the *communication graph* (abbr. *c-graph*).

Definition 2.2 (communication graph/c-graph).

A communication graph $((V, H, E), m)$ is a directed, strongly connected, labeled, bipartite graph such that:

- The graph has two kinds of nodes: *visible* nodes V and *hidden* nodes H . Each edge $e \in E$ connects two nodes of different kinds.
 - The graph has a definite root node $v_0 \in V$, each leaf is a visible node, too.
 - The labeling m maps each visible node to a set of states of the corresponding Petri net, and each edge to a bag of messages.
- *

Figure 5 shows the c-graph of the abstract BPEL process shown in Figure 2. The labeling of visible nodes was omitted, because it is not needed to verify simulation. Each edge, starting at a visible node, is labeled with a bag of messages sent by the environment – called *input*. In the root node $a1$, the module is able to consume only the message *order*. Each edge, starting at a hidden node, is labeled with a bag of messages sent by

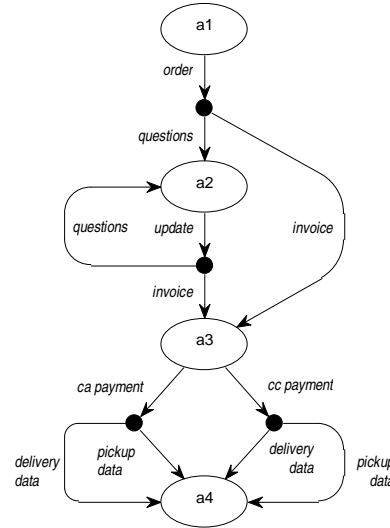


Figure 5. C-Graph of the abstract process

the module – called *output*. In the example, the abstract process answers the order by sending either the questions message or the invoice. Each path from the root to a leaf represents a complete communication sequence between module and environment. In general, an edge may be labeled with more than one message, but this is not present in the chosen examples. The c-graph of a given workflow module is well defined, and it can be calculated based on the following notions:

Activated input: Referring to a given state, an activated input is a bag of input messages that are consumed by the module along a firing sequence, whereas no output message was produced and the firing sequence ends either in the module's final state or in a state that enables a transition that can produce an output message. The function INP yields the set of activated inputs.

Successor state: Referring to a given state, a successor state is a maximal reachable state w. r. t. one possible behavior (i. e. one *concurrent run* [18]) of the module. The function NXT yields the set of successor states.

Possible output: Referring to a given state, a possible output is the bag of output message that was produced while reaching one successor state. The function OUT yields the set of possible outputs.

Communication step: A four-tuple (z, i, o, z') is called communication step if z, z' are states of a module, i is an input and o is an output, and $(z' + o)$ is a successor state of $(z + i)$. $\mathcal{S}(M)$ denotes the set of all communication steps of a module M .

As shown in Figure 5, the c-graph may contain cycles. That doesn't affect the presented analysis method as long as the graph is finite, which is the case for all

reasonable workflow modules. The precise, mathematical definition of all notions mentioned above, a discussion on the complexity of the algorithm and possible optimizations, and the problem of infinite graphs is discussed in [13]. Applying these notions, the following algorithm constructs the c-graph. The algorithm starts with the root node v_0 labeled with the initial state:

1. For each state within the label of v_k calculate the set of activated inputs: $\bigcup_{z \in m(v_k)} \text{INP}(z)$.
2. For each activated input i within this set:
 - a) Add a hidden node h , add a new edge (v_k, h) with the label i .
 - b) For each state within v_k 's label calculate the set of possible outputs: $\bigcup_{z \in m(v_k)} \text{OUT}(z + i)$.
 - c) For each possible output o within this set:
 - i. Add a visible node v_{k+1} , add a new edge (h, v_{k+1}) with the label o .
 - ii. For each state $z \in m(v_k)$ and for each communication step $(z, i, o, z') \in \mathcal{S}(M)$ add z' to the label of v_{k+1} .
 - iii. If there exists a visible node v such that $m(v_{k+1}) = m(v)$ then merge v and v_{k+1} . Otherwise, goto step 1 with node v_{k+1} .

The c-graph of a module contains the maximal information an environment can derive. In general, this graph may contain further leaf nodes. But in each c-graph there is at most one leaf node, which is labeled with the defined final state of the workflow module. All other leaf nodes contain at least one state, where there are messages left or which marks a deadlock within the module. If an environment was exchanging (with module) the sequence of messages specified by the edge labels along the path towards such a leaf node, this environment would not be an utilizing environment. The elimination of all such erroneous sequences yields a (possibly empty) subgraph that can be regarded as an user manual of the module – called the *usability graph* (abbr. *u-graph*). A workflow module is usable if and only if it has a non-empty u-graph (cf. [11]). In both examples the whole c-graph is an u-graph, too.

3 Simulation

Figure 6 presents the c-graph of the executable process model (shown in Figure 3). Comparing it manually with the graph shown in Figure 5, it is easy to figure out many similarities. In the initial state of both graphs, the customer may send the message *order*, and he will receive either the message *questions* or the message *invoice*. In case of the message *questions*, the customer should send the message *update* and may receive the message *invoice*. Finally, the customer has to pay

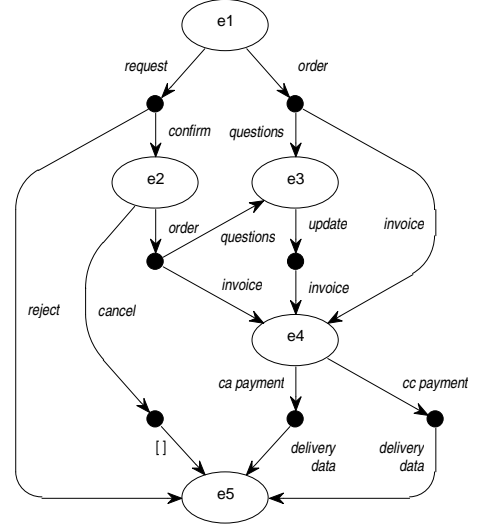


Figure 6. C-Graph of the executable process

(messages *cc payment* or *ca payment*), and the process may answer by sending the message *delivery data*.

But, there are of course many differences, too. For example, the executable process additionally accepts the message *request* at the very beginning, and it never sends the message *pickup data* to the customer. If the customer sends a *request* and cancels the process after he got the message *confirm*, the process terminates without further feedback (denoted by the empty bag []).

Hence, the executable process does not behave exactly like the abstract process. But, all differences fit a common pattern: Either a visible node has more outgoing edges than the other graph (possibly followed by structures that are not present in the other graph), or a hidden node has less outgoing edges than the other graph. Because of that, at least all those messages an utilizing environment would send to the abstract process could also be accepted by the executable process. The possible answers of the executable process are at most those, the abstract process would send to the environment. In consequence of that, such an environment can't see any difference between the two processes, i. e. the behavior of the executable process simulates the behavior of the abstract process.

3.1 Definition of simulation

To defined the simulation relation of two communication graphs precisely, some functions to navigate through those graphs are needed. Let $((V, H, E), m)$ be a c-graph of a workflow module, and let $(V_u, H_u, E_u) \subseteq (V, H, E)$ be a usability graph.

Path: A sequence of visible nodes $(v_0, \dots, v_n) \in V^*$ is called a *path* if for each two nodes v_k, v_{k+1} there is a hidden node $h \in H$, such that $(v_{k-1}, h), (h, v_k) \in E$.

Considering a path, the labels of the edges are addressed as follows: $i_k = m(v_{k-1}, h)$ and $o_k = m(h, v_k)$. *Acceptable inputs*: Given a node $v \in V_u$, an *acceptable input* is a bag of input messages that can be consumed along a path in the *usability graph*, starting at v and generating no output message in between. $AI(v)$ denotes the set of acceptable inputs in v .

Maximal paths: Given a node $v \in V_u$ and an acceptable input $i \in AI(v)$, a *maximal path* within the *c-graph* consumes at most those messages being member of i . $MP(v, i)$ denotes the set of maximal paths w.r.t. i starting at v .

Definition 3.1 (Simulation).

A communication graphs E *simulates* a communication graphs A if both graphs contain a non-empty usability graphs, and the root node of E simulates the root node of A . A node $e \in V_e$ *simulates* a node $a \in V_a$ if the following requirements are fulfilled:

- (i) For each input $i \in AI(a)$, for each path $(e, \dots, e_n) \in MP(e, i)$ there is a path $(a, \dots, a_m) \in MP(a, i)$ such that $\sum o_k^e = \sum o_k^a$ and e_n simulates a_m .
- (ii) If $AI(a)$ is empty, then $AI(e)$ is empty, too.
- (iii) For each input $i \in AI(e)$ if there is an input $i' \in AI(a)$ such that $i' \leq i$, then $i' \in AI(a)$. *

Requirement (i) has been motivated at the beginning of this section: If a module E shall simulate a module A , then E has to accept *at least* those messages A accepts, and E has to produce *at most* those messages A produces. The proper simulation of terminal states is guaranteed by requirement (ii). Finally, the module E must not respond to a smaller bag of messages than A does, e.g. sending a product without receiving the complete payment. Such behavior can yield to a deadlock, e.g. some customers might refuse to complete the payment after receiving the product. Besides the intuitive matching with our requirements the following theorem correlates the simulation of communication graphs and the simulation of workflow modules.

Theorem 3.1 (Simulation).

A workflow module E simulates a workflow module A if and only if the *c-graph* E simulates the *c-graph* A . *
As a result, *simulation* of workflow modules, i.e. *consistency* of BPEL process models can be decided effectively. The entire proof, a detailed discussion, and precise definition of all notions can be found in [13].

4 Summary

In the field of cross-organizational business processes a crucial question is to decide whether a locally defined *executable* process model is *consistent* to a globally specified *abstract* process model. Based on a Petri

net semantics for BPEL4WS [20], this paper presents a method to verify consistency automatically. The method is part of a larger framework supporting the development of Web service based business processes [14].

Additional use cases

Beside the chosen scenario of this paper, the presented method is useful in several other use cases, such as:

Process publication: If an executable process should be offered to the outside as a Web service, the abstract model (to be published) has to hide all unnecessary or confidential details of the process, and still represent the same behavior. Of course, there are approaches of syntax-based abstraction that yield a less detailed process model while preserving the consistency to the original process model (e.g. [13]). But in practice, the abstract model is often designed manually to some degree. Hence, there is a need to verify consistency between the process models at the end.

Process replacement: Assume an executable process – being already part of a cross-organizational process – needs to be replaced by another process, e.g. because of efficiency. Obviously, all other participating processes should stay untouched. In that case, the new executable process has to behave like the old one. The presented method could also be applied to verify consistency between two executable process models.

Process discovery: Assume a company is searching for a partner that provides a certain service. This company may have already a clear idea of the behavior the partner should provide. Hence, the company defines an abstract process model and wants a directory service to find a partner process, which behaves like this. This approach – called *query-by-example* in the application domain of databases – also can be supported by the consistency check presented in this paper.

Related research

The current work was inspired by many other approaches, dealing with consistency and matchmaking process models. Some approaches also use Petri nets [23, 5] and/or specify the global interaction by help of *Message Sequence Charts* (MSC) [25, 7]. But, none of them presents such a focussed view on a components externally visible behavior as the *c-graph* does. Concerning matchmaking of Web services, a recently published approach employs finite state automata to solve a similar problem [28]. It seems to yield similar theoretical results, whereas the application to a real world modeling language is not provided, yet.

Finally, there are many papers dealing with ontology and semantic Web based approaches to compare Web

services and business processes (e. g. [9]). While those methods try to match different notion with the same meaning, or to filter candidates out of a huge number of available services, the presented paper compares two given process models while having a clear idea of all used notions. Hence, semantic Web based approaches may supplement the presented method or vice versa.

Open questions

Currently, the simulation between process models is based on the assumption that there is one abstract process model for one executable process. But, an executable process modeled by help of BPEL4WS may interact with several partners, and therefore it is reasonable to have several abstract models of the same process, whereby each abstract model only emphasizes those interactions, one specific partner has to perform. In the next step, the notion of the c-graph and the definition of simulation will be adopted to meet the requirement of this scenario.

All presented algorithms are implemented within the prototype WOMBAT4WS [14]. Currently, the work is focussed on improving the algorithms' efficiency by the application of *partial order reduction* techniques [21, 19]. Moreover, up to a certain degree the integration of data aspects into the formalism is planned. Especially the dependencies between the content of incoming message and internal decisions made by the process are the focussed target. Applying technologies of static program analysis (e. g. *slicing* [15]), it seems possible, to achieve a higher level of precision in mapping a given process model into a Petri net, without loosing the possibility of efficient analysis.

References

- [1] G. Alonso, F. Casati, H. Kuno, and V. Machiraju. *Web Services*. Springer-Verlag, Dec. 2002.
- [2] Andrews et al. *BPEL- Business Process Execution Language for Web Services*. OASIS, Standard proposal, Version 1.1, July 2002.
- [3] K. Gottschalk. *Web Services architecture overview*. IBM developerWorks, Whitepaper, Sept. 2000.
- [4] P. W. Grefen and S. Angelov. *Three-Level Process Specification for Dynamic Service Outsourcing*. LNCS 2472. Springer-Verlag, 2003.
- [5] R. Hamadi and B. Benatallah. *A Petri Net based Model for Web Service Composition*. In *Proc. of ADC 2003*. Australian Computer Society, Inc., 2003.
- [6] Kavantzias, Burdett, Ritzinger, and Lafon. *Web Services Choreography Description Language*. W3C Working Draft, Version 1.0, Oct. 2004.
- [7] E. Kindler, A. Martens, and W. Reisig. *Inter-operability of Workshop Applications - Local Criteria for Global Soundness*. LNCS 1806. Springer-Verlag, 2000.
- [8] F. Leymann and D. Roller. *Modeling Business Processes with BPEL4WS*. In *Proc. of Modellierung 2004*, LNI 45. GI, 2004.
- [9] S. A. Ludwig and P. van Santen. *A Grid Service Discovery Matchmaker Based on Ontology Description*. In *EuroWeb 2002*. British Computer Society, 2002.
- [10] A. Martens. On compatibility of web services. *Petri Net Newsletter*, 65:12-20, Oct. 2003.
- [11] A. Martens. *On Usability of Web Services*. In Calero, Diaz, and Piattini, editors, *Proceedings of WQW 2003*, Rome, Italy, Dec. 2003. IEEE Computer Society Press.
- [12] A. Martens. *Analysis and Re-engineering of Web Services*. In S. Hammoudi and J. Cordeiro, editors, *Proceedings of ICEIS 2004*, Porto, Portugal, 2004.
- [13] A. Martens. *Verteilte Geschäftsprozesse - Modellierung und Verifikation mit Hilfe von Web Services*. PhD thesis, WiKu-Verlag Stuttgart, 2004.
- [14] A. Martens et al. *BPEL - Semantik, Analyse und Visualisierung*. Informatik-Bericht 169, Humboldt-Universität zu Berlin, 2004.
- [15] F. Nielson, H. R. Nielson, and C. Hankin. *Principles of Program Analysis*. Springer-Verlag, 1999.
- [16] A. Norta. *Web Supported Enactment of Petri-Net Based Workflows with XRL/flower*. LNCS 3099. Springer-Verlag, June 2004.
- [17] W. Reisig. *Petri Nets*. Springer-Verlag, 1985.
- [18] W. Reisig. *Elements of Distributed Algorithms - Modeling and Analysis with Petri Nets*. Springer, 1998.
- [19] K. Schmidt. *Explicit State Space Verification*. Habilitation, Humboldt-Universität zu Berlin, 2002.
- [20] K. Schmidt and C. Stahl. A Petri net semantic for BPEL. In E. Kindler, editor, *Proc. of 11th Workshop AWPN*. Paderborn University, Oct. 2004.
- [21] A. Valmari. *State Space Generation - Efficiency and Practicality*. PhD thesis, Tampere University, 1988.
- [22] W. van der Aalst and A. ter Hofstede. *YAWL: Yet Another Workflow Language*. Brisbane, Australia, 2003.
- [23] W. M. P. van der Aalst. *Modeling and Analyzing Interorganizational Workflows*. In *Proc. of CSD'98*. IEEE Computer Society Press, 1998.
- [24] W. M. P. van der Aalst. The Application of Petri Nets to Workflow Management. *Journal of Circuits, Systems and Computers*, 8(1):21-66, 1998.
- [25] W. M. P. van der Aalst. Interorganizational Workflows - An Approach based on MSC and Petri Nets. *Systems Analysis - Modelling - Simulation*, 34(3), 1999.
- [26] R. J. van Glabbeek. *The Linear Time - Branching Time Spectrum*. LNCS 458. Springer-Verlag, 1990.
- [27] Web-Ontology Working Group. *Ontology Web Language for Web Services*. OWL-S 1.1, Nov. 2004.
- [28] A. Wombacher, P. Fankhauser, B. Mahleko, and E. Neuhold. *Matchmaking for Business Processes*. In *Proc. of EEE-04*. IEEE Computer Society, 2004.
- [29] WOMBAT4WS. *Workflow Modeling and Business Analysis Toolkit for Web Services*. homepage, <http://www.informatik.hu-berlin.de/top/wombat/>.