

AIOCJ: A Choreographic Framework for Safe Adaptive Distributed Applications

Technical Report

Mila Dalla Preda¹, Saverio Giallorenzo²,
Ivan Lanese², Jacopo Mauro², and Maurizio Gabbriellini²

¹ Department of Computer Science - Univ. of Verona

² Department of Computer Science and Engineering - Univ. of Bologna / INRIA

Abstract. We present AIOCJ, a framework for programming distributed adaptive applications. Applications are programmed using AIOC, a choreographic language suited for expressing patterns of interaction from a global point of view. AIOC allows the programmer to specify which parts of the application can be adapted. Adaptation takes place at runtime by means of rules, which can change during the execution to tackle possibly unforeseen adaptation needs. AIOCJ relies on a solid theory that ensures applications to be deadlock-free by construction also after adaptation. We describe the architecture of AIOCJ, the design of the AIOC language, and an empirical validation of the framework.

1 Introduction

Adaptation is a main feature of current distributed applications, that should live for a long time in a continuously changing environment. Anticipating all the possible adaptation needs when designing an application is very difficult, thus the approaches able to cope with *unforeseen* adaptation needs are the most interesting. Also, for distributed applications like the ones that we consider, it is important to ensure deadlock-freedom (according to [1] about one third of concurrency bugs in real applications are deadlocks). While many techniques ensuring deadlock freedom exist in the literature, e.g., [2–4], to the best of our knowledge, none of them deals with adaptive applications. Indeed, most of the approaches to adaptation offer no guarantee on the behaviour of the application after adaptation [5–7], or they assume to know all the possible adaptations in advance [8], thus failing to cope with unforeseen adaptation needs.

Here we present AIOCJ, a prototype implementation of a framework for programming adaptive distributed applications that guarantees deadlock-freedom by construction (the theoretical foundations ensuring this property are discussed in [9]). AIOCJ is composed of two parts: (i) a domain-specific language, called Adaptive Interaction-Oriented Choreographies (AIOC) and (ii) an *adaptation middleware* that supports adaptation of AIOC programs.

The AIOC language describes applications from a global point of view following the *choreography* paradigm. This paradigm has been applied in different contexts, see, e.g., [2, 10–13], but we are not aware of other tools based on it and targeting adaptive applications. A choreography defines the interactions among the processes of a

distributed application. AIOC main innovation consists in two constructs supporting adaptation: *scopes* and *adaptation rules*. A scope delimits code that may be adapted in the future. An adaptation rule provides new code to replace the one in a given scope. Interestingly, in AIOCJ, adaptation rules can be defined and inserted in the framework while the application is running, to cope with adaptation needs which were not foreseen when the application was designed or even started.

The code below shows a toy AIOC program (left) and an adaptation rule applicable to it (right). On the left, Lines 2-4 define a scope in which the local variable `msg` of process `user` is set to `"Hello World"`. The keyword `prop` defines properties of the scope (prefixed by `N`). In this case, the name property is set to `"hello_world"`. At Line 5 `user` sends the content of `msg` to a second process (`display`), that stores it in its local variable `msg`. On the right, Lines 2-3 define the *applicability condition* of the rule, i.e., the name property of the scope should be set to `"hello_world"`, and the environmental property `E.lang` should be equal to `"it"`. Line 4 shows the code that will replace the one of the scope, i.e., variable `msg` of `user` will be set to `"Ciao Mondo"` (Italian for "Hello World").

```

1  aioc {
2    scope @user{
3      msg@user = "Hello World"
4    } prop { N.name = "hello_world" };
5    send: user( msg ) -> display( msg ) }

1  rule {
2    on { N.name == "hello_world"
3      and E.lang == "it" }
4    do { msg@user = "Ciao Mondo" }
5  }

```

An AIOC program describes a full distributed application. AIOCJ generates Jolie code for each (possibly distributed) process. Jolie [14, 15] is a Service-Oriented orchestration language, and we chose it since it provides primitives suitable for implementing our communication and adaptation behaviour.

The adaptation middleware consists of a set of adaptation rules stored in multiple, possibly distributed, *adaptation servers*, and of an *adaptation manager* that mediates the interactions between the adaptive application and the various adaptation servers.

Structure of the paper. Section 2 presents an overview of the AIOCJ framework, while Section 3 describes its implementation. Section 4 shows a preliminary validation of the framework, with tests on the performances of AIOCJ. In Section 5 we discuss related work and future directions of research. Appendix A outlines a short demo of the use of the framework.

2 Overview: The AIOCJ framework

This section first defines the architectural model that supports adaptation of AIOCJ applications and then it introduces the syntax of the AIOC language via an example (for a formal presentation of AIOC syntax and semantics see [9]).

The AIOCJ middleware. We consider applications composed of processes deployed as services on different localities, including local state and computational resources. Each process has a specific duty in the choreography and follows a given protocol. Processes interact via synchronous message passing over channels, also called *operations*. Adaptation is performed by an adaptation middleware including an adaptation manager and some, possibly distributed, adaptation servers. The latter are services that act as repositories of adaptation rules and may be (manually) added or removed at runtime. Running

adaptation servers register themselves on the adaptation manager. The running application may interact with the adaptation manager to look for applicable adaptation rules. The effect of an adaptation rule is to replace a scope with new code that answers a given adaptation need. The adaptation manager checks the rules according to the registration order of the adaptation servers, returning the first applicable rule, if any.

The AIOC Language. The language relies on a set of roles that identify the processes in the choreography. Let us introduce the syntax of the language using an example where Bob invites Alice to see a film.

```

1  include isFreeDay from "calendar.org:80" with http
2  include getTicket from "cinema.org:8000" with soap
3  preamble {
4    starter: bob
5    location@bob = "socket://localhost:8000"
6    location@alice = "socket://alice.com:8000"
7    location@cinema = "socket://cinema.org:8001" }
8  aioc{
9    end@bob = false;
10   while( ! end )@bob{
11     scope @bob {
12       free_day@bob = getInput( "Insert your free day" );
13       proposal: bob( free_day ) -> alice( bob_free_day );
14       is_free@alice = isFreeDay( bob_free_day );
15     } prop { N.scope_name = "matching day" };
16     if( is_free )@alice {
17       scope @bob {
18         proposal: bob( "cinema" ) -> alice( event );
19         agreement@alice = getInput( "Bob proposes " + event +
20           ", do you agree?[y/n]");
21         if( agreement == "y" )@alice{
22           end@bob = true;
23           book: bob( bob_free_day ) -> cinema( book_day );
24           ticket@cinema = getTicket( book_day );
25           { notify: cinema( ticket ) -> bob( ticket )
26             | notify: cinema( ticket ) -> alice( ticket ) }}
27         } prop { N.scope_name = "event selection" };
28       } if( !end )@bob {
29         _r@bob = getInput( "Alice refused. Try another date?[y/n]" );
30         if( _r != "y" )@bob{ end@bob = true }}}

```

Listing 1.1: Appointment program.

The code starts with some deployment information (Lines 1-7), discussed later on. The behaviour starts at Line 9. The program is made by a cycle where Bob first checks when Alice is available and then invites her to the cinema. Before starting the cycle, Bob initialises the variable `end`, used in the guard of the cycle, to the boolean value `false` (Line 9). Note the annotation `@bob` meaning that `end` is local to Bob. The first instructions of the while are enclosed in a scope (Lines 11-15), meaning that they may be adapted in the future. The first operation within the scope is the call to the primitive function `getInput` that asks to Bob a day where he is free and stores this date into the local variable `free_day`. At Line 13 the content of `free_day` is sent to Alice via operation `proposal`. Alice stores it in its local variable `bob_free_day`. Then, at Line 14, Alice calls the external function `isFreeDay` that checks whether she is available on `bob_free_day`. If she is available (Line 16) then Bob sends to her the invitation to go to the cinema via the operation `proposal` (Line 18). Alice, reading from the input, accepts or refuses the invitation (Line 19). If Alice accepts then Bob first sets the variable `end` to `true` to end the cycle. Then, he sends to the cinema the booking re-

quest via operation book. The cinema generates the tickets using the external function `getTicket` and sends them to Alice and Bob via operation `notify`. The two notifications are done in parallel using the parallel operator `|` (until now we composed statements using the sequential operator `;`). Lines 18-26 are enclosed in a second scope with property `N.scope_name = "event selection"`. If the agreement is not reached, Bob decides, reading from the input, if he wants to stop inviting Alice. If so, the program exits.

We remark the different meanings of the annotations `@bob` and `@alice`. When prefixed by a variable, they identify the owner of the variable. Prefixed by the boolean guard of conditionals and cycles, they identify the role that evaluates the guard. Prefixed by the keyword `scope`, they identify the process coordinating the adaptation of that scope. A scope, besides the code, may also include some properties describing the current implementation. These can be specified using the keyword `prop` and are prefixed by `N`. For instance, each scope of the example includes the property `scope_name`, that can be used to distinguish its functionality.

AIOGJ can interact with external services, seen as functions. This allows both to interact with real services and to have easy access to libraries from other languages. To do that, one must specify the address and protocol used to interact with them. For instance, the external function `isFreeDay` used in Line 14 is associated to the service deployed at the domain “calendar.org”, reachable through port 80, and that uses `http` as serialisation protocol (Line 1). External functions are declared with the keyword `include`. To preserve deadlock freedom, external services must be non-blocking. After function declaration, in a `preamble` section, it is possible to declare the locations where processes are deployed. The keyword `starter` is mandatory and defines which process must be started first. The starter makes sure all other processes are ready before the execution of the choreography begins.

Now suppose that Bob, during summer, prefers to invite Alice to a picnic more than to the cinema, provided that the weather forecasts are good. This can be obtained by adding the following adaptation rule to one of the adaptation servers. This may even be done while the application is running, e.g., while Bob is sending an invitation. In this case, if the first try of Bob is unsuccessful, in the second try he will propose a picnic.

```

1 rule {
2   include getWeather from "socket://localhost:8002"
3   on { N.scope_name == "event selection" and E.month > 5 and E.month < 10 }
4   do { forecasts@bob = getWeather( free_day );
5     if( forecasts == "Clear" )@bob{
6       eventProposal: bob( "picnic" ) -> alice( event )
7     } else { eventProposal: bob( "cinema" ) -> alice( event ) };
8   agreement@alice = getInput( "Bob proposes " + event +
9     ", do you agree?[y/n]");
10  if( agreement == "y" )@alice {
11    end@bob = true |
12    if( event == "cinema" )@alice {
13      //cinema tickets purchase procedure
14  }}}

```

Listing 1.2: Event selection adaptation rule.

A rule specifies its applicability condition and the new code to execute. In general, the applicability condition may depend only on properties of the scope, environment variables, and variables belonging to the coordinator of the scope. In this case, the condition, introduced by the keyword `on` (Line 3), makes the rule applicable to scopes

having the property `scope_name` equal to the string `"event selection"` and only during summer. This last check relies on an environment variable `month` that contains the current month. Remarkably, one can argue that the rule is correct only in the northern hemisphere. However, when the application runs in the southern hemisphere, a different adaptation server will be available, containing a rule tailored for the southern hemisphere. Environment variables are prefixed by `E`.

When the rule applies, the new code to execute is defined using the keyword `do` (Line 4). In this case, the forecasts can be retrieved calling an external function `getWeather` (Line 4) that queries a weather forecasts service. This function is declared in Line 2. If the weather is clear, Bob proposes to Alice a picnic, the cinema otherwise. Booking (as in Listing 1.1, Lines 23-26) is needed only if Alice accepts the cinema proposal.

As detailed in [9], to obtain a deadlock-free application, we require the code of choreographies and rules to satisfy a well-formedness syntactic condition called *connectedness*. Intuitively, connectedness ensures that sequences of actions are executed in the correct order and avoids interference between parallel interactions. Requiring this condition does not hamper programmability, since it naturally holds in most of the cases, and it can always be enforced automatically via small patches to the choreography which preserve the behaviour of the program, as discussed in [16]. Also, checking connectedness is efficient, i.e., polynomial w.r.t. the size of the code [9].

3 Implementation

Our prototype implementation of AIOCJ is composed of two elements: the AIOCJ Integrated Development Environment (IDE), named `AIOCJ-ecl`, and the adaptation middleware that enables AIOC programs to adapt, called `AIOCJ-mid`.

`AIOCJ-ecl` is a plug-in for Eclipse [17] based on Xtext [18]. Xtext provides features such as syntax highlighting, syntax checking, and code completion, which help developers in writing choreographies and adaptation rules. Also, starting from a grammar, Xtext generates the parser for programs written in the AIOC language. Result of the parsing is an abstract syntax tree (AST) we use to implement (i) the checker for connectedness for choreographies and rules and (ii) the generation of Jolie code for each role. The connectedness check has polynomial computational complexity [9] thus making it efficient enough to be performed on-the-fly while editing the code.

The target language of code generation is Jolie [14]. Jolie supports architectural primitives such as dynamic embedding, aggregation, and redirection that we exploit to implement the adaptation mechanisms. Moreover, Jolie supports a wide range of communication technologies (TCP/IP sockets, local memory, Bluetooth) and of data formats (e.g., HTTP, SOAP, JSON). AIOCJ inherits this ability. Actually, any language that supports runtime code evaluation can implement our adaptation mechanisms. However, we deem that using Jolie architectural primitives particularly suits the requirements of our framework. The compilation generates a Jolie service for each role. The execution of scopes is delegated to sub-services accessed using Jolie redirection facility. Adaptation is enacted by disabling the current sub-service and replacing it with a new one, obtained from the adaptation server. To grant to all the sub-services access to variables, the state is stored by a dedicated sub-service local to the role. Auxiliary messages are

exchanged to ensure that both the adaptation and the choices taken by the if and while constructs are done in a coordinated way. In particular, the scope execution not only requires interaction with the adaptation manager, but also communications among the different roles, ensuring that they all agree on whether adaptation is needed or not, and, in case, on which rule to apply. Indeed, the decision is taken by the role coordinating the adaptation and then communicated to other roles. Note that the different roles cannot autonomously take the decision, since if they take it at different times, changes in the environment or in the sets of available rules may lead to inconsistent decisions.

Synchronous message exchange is implemented on top of an asynchronous communication middleware by a sub-service that works as a *message handler*. The message handler of the starter role also ensures that, before the actual communication in the choreography starts, all the roles are ready.

AIOCJ-mid is implemented in Jolie and it includes:

- many, possibly distributed, *adaptation servers* where rules are published. Adaptation servers can be deployed and switched on and off at runtime;
- an *adaptation manager* that acts as a registry for adaptation servers and clients;
- an *environment* service that stores and makes available environment information. Environment information can change at any moment.

When an AIOCJ program reaches a scope, it queries the adaptation manager for a rule matching that scope. The adaptation manager queries each adaptation server sequentially, based on their order of registration. Each server checks the applicability condition of each of its rules. The first rule whose applicability condition holds is applied. In particular, the code of the rule is sent to the role coordinating the adaptation (via the adaptation manager) which distributes it to the involved roles. In each role, the new code replaces the old one. The study of more refined policies for rule selection, e.g., based on priorities, is a topic for future work.

4 Validation

In this section, we give a preliminary empirical validation of our implementation. The main aim is to test how our mechanisms for adaptation impact on performances.

In the literature, to the best of our knowledge, there is no approach to adaptation based on choreography programming. Thus, it is difficult to directly compare our results with other existing approaches. Moreover, we are not aware of any established benchmark to evaluate adaptive applications. For this reason, we tested AIOCJ performances by applying it to two typical programming patterns: *pipes* and *fork-joins*. Since we are interested in studying the cost of adaptation, our scenarios contain minimal computation and are particularly affected by the overhead of the adaptation process. Clearly, the percentage of the overhead due to adaptation will be far lower in real scenarios, which are usually more computationally intensive. In the first scenario, we program a pipe executing n tasks (in a pipe, the output of task t_i is given as input to task t_{i+1} , for $i \in \{1, \dots, n - 1\}$). To keep computation to a minimum, each task simply computes the increment function. In the fork-join scenario, n tasks are computed in parallel. Each

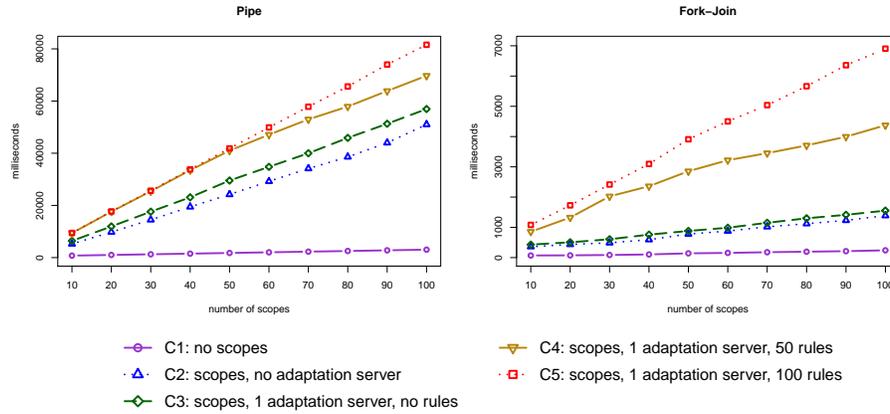


Fig. 1: Times of execution of the pipe (left) and the fork-join (right) scenarios

task processes one character of a message of length n , shifting it by one position. The message is stored in an external service.³

To enable adaptation, each task is enclosed in a scope. We test both scenarios with an increasing number of tasks $n \in \{10, 20, \dots, 100\}$ to study how performances scale as the number of adaptation scopes increases. We evaluate performances in different contexts, thus allowing us to understand the impact of different adaptation features, such as scopes, adaptation servers, and adaptation rules.

Context 1: no scopes, no adaptation servers, no rules;

Context 2: each task is enclosed in a scope, no adaptation servers, no rules;

Context 3: each task is enclosed in a scope, one adaptation server, no rules;

Context 4: as Context 3, but now the adaptation server contains 50 rules. Each rule is applicable to a unique scope i , and no rule is applicable to scopes with $i > 50$. The rules are stored in random order.

Context 5: as Context 4, but with 100 rules, one for each scope.

Each rule in Contexts 4 and 5 is applicable to one specific scope only (through a unique property of the scope), hence when testing for 50 rules, only the first 50 scopes adapt.

In order to reduce variability of performances, we repeated every test 5 times. We performed our tests on a machine equipped with a 2.6GHz quad-core Intel Core i7 processor and 16GB RAM. The machine runs Mavericks 10.9.3, Java 1.7.55, and Jolie r.2728. Figure 1 shows the tests for the pipe (left) and the fork-join (right). Both charts display on the x-axis the number of tasks/scopes and on the y-axis the execution time in milliseconds.

As expected, in both scenarios there is a significant gap between Contexts 1 and 2. In words, the introduction of scopes has a strong effect on performances. The ratio is 1:13 for the pipe scenario and 1:5.5 for the fork-join scenario. This is due to the auxiliary communications needed to correctly execute a scope. The observed overhead

³ The code of both scenarios is in Appendix B.

is higher in the pipe scenario, since different scopes check for adaptation in sequence, while this is done in parallel for the fork-join scenario.

Adding an adaptation server (from Context 2 to Context 3) has little impact on performances: 19% of decay for pipe, and 17% for fork-join. The figures are reasonable, considered that Context 3 adds only one communication w.r.t. Context 2.

On the contrary, there is a notable difference when adding rules to the adaptation server (Context 4 is 1.4 times slower than Context 3 for the pipe scenario, 2.9 for the fork-join scenario). In Contexts 4 and 5, performances are really close up to 50 scopes (in the pipe scenario they almost overlap) although Context 5 has twice the rules of Context 4. This illustrates that the time to test for applicability of rules is negligible. Hence, the highest toll on performances is related to actual adaptation, since it requires to transfer and embed the new code. This is particularly evident in the fork-join scenario where multiple adaptations are executed in parallel and the adaptation server becomes a bottleneck. This problem can be mitigated using multiple distributed adaptation servers.

The fact that the most expensive operations are scope execution and actual adaptation is highlighted also by the results below. The table shows the cost of different primitives, including scopes in different contexts. Times are referred to 5 executions of the sample code in Appendix B.

Test	Time (ms)	Test	Time (ms)
assignment	2.2	scope, 1 adaptation server, 1 matching rule	280.6
interaction	4.2	scope, 1 adaptation server, 50 rules, none matching	254.2
if statement	16.6	scope, 1 adaptation server, 50 rules, 1 matching	338.6
scope, no adaptation server	129.4	scope, 1 adaptation server, 100 rules, none matching	310.2
scope, 1 adaptation server, no rule	203.8	scope, 1 adaptation server, 100 rules, 1 matching	385

As future work we will exploit these results to increase the performances of our framework, concentrating on the bottlenecks highlighted above. For instance, scope execution (as well as conditionals and cycles) currently requires many auxiliary communications ensuring that all the processes agree on the chosen path. In many cases, some of these communications are not needed, since a process will eventually discover the chosen path from the protocol communications. Static analysis can discover redundant communications and remove them. Another improvement is letting the adaptation server send the new code directly to the involved roles, skipping the current forward chain.

5 Related Work & Conclusion

This paper presented a framework for programming rule-based adaptation of distributed applications. Its distinctive trait is that, being based on a choreographic approach, it guarantees deadlock-freedom by construction for the running distributed application, even in presence of adaptation rules which were unknown when the application was started, and for any environment condition.

Adaptation is a hot topic, and indeed there is a plethora of approaches in the literature, see, e.g., the surveys [19, 20]. However, approaches based on formal methods are only emerging recently and few of them have been implemented in a working tool. In particular, the use of choreographies to capture and define adaptive applications is a novel idea. For a discussion of works on adaptation with formal bases, but which have not been implemented, we refer to [9]. Here, we just recall [21], which exploits a choreographic approach for self-adaptive monitoring of distributed applications.

Among the implemented approaches, the most related to ours is JoRBA [5]. JoRBA features scopes and adaptation rules similar to ours. However, JoRBA applications are not distributed and JoRBA does not guarantee any property of the adapted application.

In [22] choreographies are used to propagate protocol changes to the other peers, while [23] presents a test to check whether a set of peers obtained from a choreography can be reconfigured to match a second one. Differently from ours, these works only provide change recommendations for adding and removing message sequences.

Various tools [24–26] exploit automatic planning techniques in order to elaborate, at runtime, the best sequence of activities to achieve a given goal. These techniques are more declarative than ours, but, to the best of our knowledge, they are not guaranteed to always find a plan to adapt the application.

Among the non-adaptive languages, Chor [2] is the closest to ours. Indeed, like ours, Chor is a choreographic language that compiles to Jolie. Actually, AIOCJ shares part of the Chor code base. However, due to the different semantics of the sequential operator and the lack of the parallel composition in Chor, a faithful encoding of the scenarios in Section 4 is not possible, especially for the fork-join scenario. On an almost equivalent implementation of the pipe scenario, Chor proves to be more efficient than AIOCJ. This performance gap reinforces our idea that the current implementation of AIOCJ can be optimised, possibly taking inspiration from some implementation choices made in Chor.

In the future, we would like to test the expressive power of our language, trying to encode patterns of adaptation from existing approaches. An obvious benefit of such an encoding is that it will capture patterns of adaptation used in real-case scenarios, guaranteeing also deadlock freedom, which is not provided by other approaches. This task is cumbersome, due to the huge number and heterogeneity of those approaches. Nevertheless, we already started it. In particular, in the website [27], we show how to encode examples coming from distributed [28] and dynamic [29] Aspect-Oriented Programming (AOP) and from Context-Oriented Programming (COP) [30]. In general, we can deal with cross-cutting concerns like logging and authentication, typical of AOP, viewing point-cuts as empty scopes and advices as adaptation rules. Layers, typical of COP, can instead be defined by adaptation rules which can fire according to contextual conditions captured by the environment. Possible extensions of our framework include the use of asynchronous communications in the AIOC language and the introduction of mechanisms to deal with exceptions and failures. Finally, we would like to pursue a systematic analysis of the workflow change patterns like the ones presented in [31, 32], showing how these patterns are captured by AIOCJ.

References

1. S. Lu, S. Park, E. Seo, and Y. Zhou, "Learning from mistakes: a comprehensive study on real world concurrency bug characteristics," in *ASPLOS*, pp. 329–339, ACM, 2008.
2. M. Carbone and F. Montesi, "Deadlock-Freedom-by-Design: Multiparty Asynchronous Global Programming," in *POPL*, pp. 263–274, ACM, 2013.
3. G. Göbller and J. Sifakis, "Component-Based Construction of Deadlock-Free Systems," in *FSTTCS*, vol. 2914 of *LNCS*, pp. 420–433, Springer, 2003.
4. M. Naik, C.-S. Park, K. Sen, and D. Gay, "Effective static deadlock detection," in *ICSE*, pp. 386–396, IEEE, 2009.
5. I. Lanese, A. Bucchiarone, and F. Montesi, "A Framework for Rule-Based Dynamic Adaptation," in *TGC*, vol. 6084 of *LNCS*, pp. 284–300, Springer, 2010.
6. A. Bucchiarone, A. Marconi, M. Pistore, and H. Raik, "Dynamic Adaptation of Fragment-Based and Context-Aware Business Processes," in *ICWS*, pp. 33–41, IEEE Press, 2012.
7. W.-K. Chen, M. A. Hiltunen, and R. D. Schlichting, "Constructing Adaptive Software in Distributed Systems," in *ICDCS*, vol. 6084 of *LNCS*, pp. 635–643, Springer, 2001.
8. J. Zhang, H. Goldsby, and B. H. C. Cheng, "Modular Verification of Dynamically Adaptive Systems," in *AOSD*, pp. 161–172, ACM, 2009.
9. M. Dalla Preda, M. Gabbriellini, S. Giallorenzo, I. Lanese, and J. Mauro, "Deadlock Freedom by Construction for Distributed Adaptive Applications." Technical Report. <http://arxiv.org/pdf/1407.0970v1.pdf>.
10. M. Carbone, K. Honda, and N. Yoshida, "Structured communication-centered programming for web services," *ACM Trans. Program. Lang. Syst.*, vol. 34, no. 2, p. 8, 2012.
11. "Scribble website." <http://www.jboss.org/scribble>.
12. I. Lanese, C. Guidi, F. Montesi, and G. Zavattaro, "Bridging the Gap between Interaction- and Process-Oriented Choreographies," in *SEFM*, pp. 323–332, IEEE Press, 2008.
13. World Wide Web Consortium, *Web Services Choreography Description Language Version 1.0*, 2005. <http://www.w3.org/TR/ws-cdl-10/>.
14. "Jolie website." <http://www.jolie-lang.org/>.
15. F. Montesi, C. Guidi, and G. Zavattaro, "Composing services with JOLIE," in *Proc. of ECOWS'07*, pp. 13–22, IEEE Press, 2007.
16. I. Lanese, F. Montesi, and G. Zavattaro, "Amending choreographies," in *WWV*, vol. 123, pp. 34–48, EPTCS, 2013.
17. "Eclipse website." <http://www.eclipse.org/>.
18. "Xtext website." <http://www.eclipse.org/Xtext/>.
19. C. Ghezzi, M. Pradella, and G. Salvaneschi, "An evaluation of the adaptation capabilities in programming languages," in *SEAMS*, pp. 50–59, ACM, 2011.
20. L. A. F. Leite *et al.*, "A systematic literature review of service choreography adaptation," *Service Oriented Computing and Applications*, vol. 7, no. 3, pp. 199–216, 2013.
21. M. Coppo, M. Dezani-Ciancaglini, and B. Venneri, "Self-adaptive monitors for multiparty sessions," in *PDP*, pp. 688–696, IEEE, 2014.
22. S. Rinderle, A. Wombacher, and M. Reichert, "Evolution of Process Choreographies in DY-CHOR," in *OTM Conferences (1)*, vol. 4275 of *LNCS*, pp. 273–290, Springer, 2006.
23. A. Wombacher, "Alignment of choreography changes in BPEL processes," in *IEEE SCC*, pp. 1–8, IEEE Press, 2009.
24. G. Cugola, C. Ghezzi, and L. S. Pinto, "DSOL: a declarative approach to self-adaptive service orchestrations," *Computing*, vol. 94, no. 7, pp. 579–617, 2012.
25. L. Baresi, A. Marconi, M. Pistore, and A. Sirbu, "Corrective Evolution of Adaptable Process Models," in *BMMDS/EMMSAD*, vol. 147 of *LNBIIP*, pp. 214–229, Springer, 2013.

26. A. Bucchiarone, A. Marconi, C. A. Mezzina, M. Pistore, and H. Raik, "On-the-Fly Adaptation of Dynamic Service-Based Systems: Incrementality, Reduction and Reuse," in *ICSOC*, vol. 8274 of *LNCS*, pp. 146–161, Springer, 2013.
27. "AIOCJ website." <http://www.cs.unibo.it/projects/jolie/aiocj.html>.
28. R. Pawlak *et al.*, "JAC: an aspect-based distributed dynamic framework," *SPE*, vol. 34, no. 12, pp. 1119–1148, 2004.
29. Z. Yang, B. H. C. Cheng, R. E. K. Stirewalt, J. Sowell, S. M. Sadjadi, and P. K. McKinley, "An aspect-oriented approach to dynamic adaptation," in *WOSS*, pp. 85–92, ACM, 2002.
30. R. Hirschfeld, P. Costanza, and O. Nierstrasz, "Context-oriented Programming," *Journal of Object Technology*, vol. 7, no. 3, pp. 125–151, 2008.
31. B. Weber, S. Rinderle, and M. Reichert, "Change Patterns and Change Support Features in Process-Aware Information Systems," in *CAiSE*, pp. 381–395, Springer, 2013.
32. F. Casati, S. Ceri, B. Pernici, and G. Pozzi, "Workflow Evolution," *Data Knowl. Eng.*, vol. 24, no. 3, pp. 211–238, 1998.

A Demonstration

In this section, we describe, supported by screenshots, how a developer can write an AIOCJ application and its adaptation rules. We also show an example of execution and adaptation of the generated application.

A.1 Writing an AIOCJ application

We consider the AIOC program showed in Section 2, but we adapt the deployment of external services and of the roles to run them locally. Note that, when not explicitly defined in the **preamble**, the location of a role is automatically assigned to a local address.

```
1 include isFreeDay from "socket://localhost:8000"
2 include getTicket from "socket://localhost:8001"
3
4 preamble { starter: bob }
5
6 aioc {
7   // Code as in Listing 1.1
8 }
```

Listing 1.3: Local deployment for appointment program.

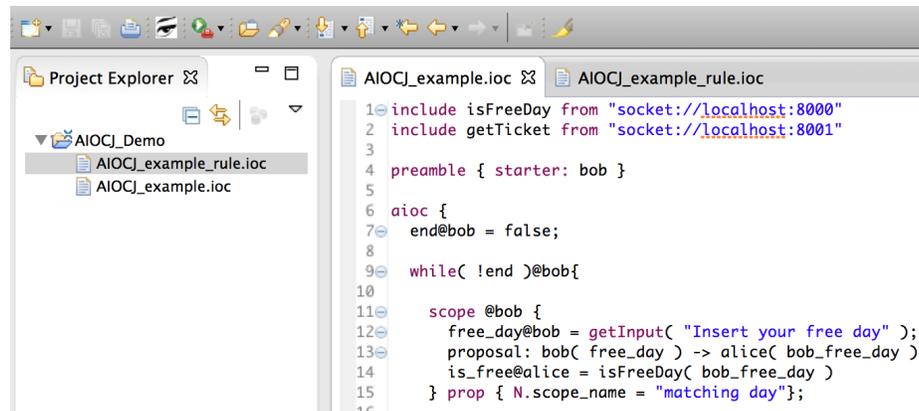


Fig. 2: The AIOCJ eclipse plug-in

Fig. 2 shows the interface of the Eclipse IDE running the AIOCJ-ecl plug-in. In AIOCJ-ecl developers can write choreographies and adaptation rules. The IDE provides syntax checking and highlighting. Plus, it performs an online check for connectivity.

A.2 Online check for connectedness

To see an example of how the online check for connectedness works, let us modify the code in Listing 1.1 to generate a non-connected sequence. We switch Lines 13 and 14 of Listing 1.1 as shown in the snippet below.

```
12 free_day@bob = getInput( "Insert your free day" );
13 is_free@alice = isFreeDay( bob_free_day );
14 proposal: bob( free_day ) -> alice( bob_free_day )
```

Listing 1.4: Non-connected sequence.

Fig. 3 shows AIOCJ-ecl notifying the error on the first non-connected instruction. Note that the code in Listing 1.4 does not behave as expected, since Alice uses value `bob_free_day` at Line 13 before receiving it at Line 14. Indeed, connectedness errors frequently highlight errors in the logic of the application.

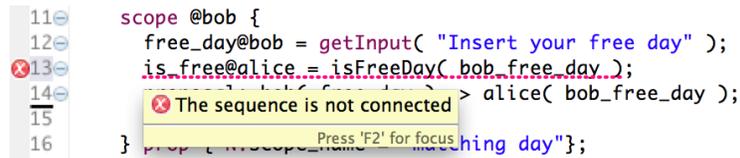


Fig. 3: Check for connectedness: sequence.

Similarly, AIOCJ-ecl also notifies when connectedness fails on a parallel operator (Fig. 4). In this case, w.r.t. the code in Listing 1.1, we miswrote the notification procedure at Lines 25-26, sending in parallel twice the same message from cinema to Bob.

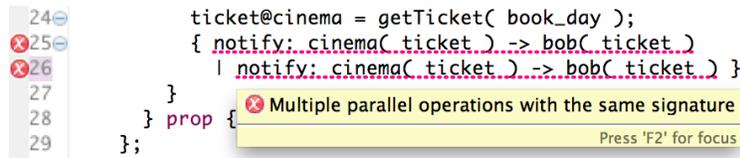


Fig. 4: Check for connectedness: parallel.

A.3 Compiling and Running an AIOCJ application

The icon in the top toolbar of AIOCJ-ecl highlighted in Fig. 5 starts the compilation of the AIOC program to Jolie. A dialog box notifies the successful compilation of the program.

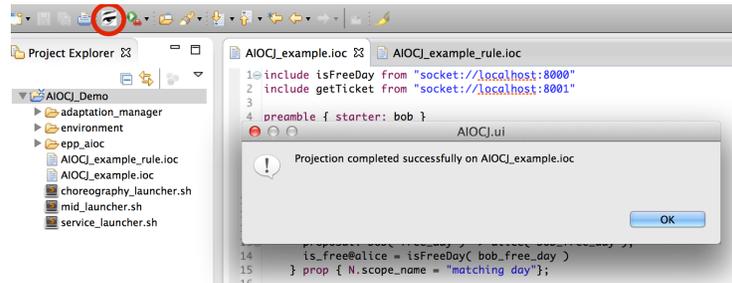


Fig. 5: Compilation of the AIOCJ application.

The compilation of an AIOC application generates three folders, as showed in the left column of Fig. 5. The folders contain the adaptation manager, the environment, and the roles in the choreography (in folder “epp_aioc”).

AIOCJ-ecl also projects utilities for testing the choreography locally. The script “choreography_launcher.sh” takes care of launching, in the proper order, the processes of the adaptation middleware (the environment, the adaptation manager, and an adaptation server, if present). Then the script launches the roles in the choreography, beginning with the **starter** one.

Listing 1.5 shows as an example the Jolie code generated from the choreography in Listing 1.1 for Bob. We highlight below how some of the primitives of the AIOC language are encoded in Jolie.

Jolie provides **outputPorts** for defining how services communicate. In our case, Bob has one **outputPort** for each process he communicates with (i.e., alice and cinema). Output ports are defined at Lines 1-11.

Lines 18-21 encode the assignment `end@bob = false`. At Line 18 the variable `end` is set to **false**. Then, the value is sent to the State sub-service which is in charge of storing it (Lines 19-21). Notice that this code is needed since `end` is a local variable of Bob. In contrast, the assignment `is_free@alice = isFreeDay(bob_free_day)` is done by Alice and it does not appear in Bob’s code.

Lines 22-30 show the encoding of the **while** (Listing 1.1, Line 10). First the variable `end` is retrieved from the State sub-service (Line 22). Then the guard `!end` is evaluated (Line 23), and the result of the evaluation is sent to the roles involved in the cycle (Lines 24-29). Finally, at Line 30, Bob enters the cycle. Notably, the sequence at Lines 22-29 is repeated at Lines 58-65 to update the evaluation of the guard of the cycle.

At Line 31, Bob enters the scope named “**matching day**”. The code of the scope is not reported here since scopes are externalised as auxiliary sub-services. However, one can see the invocation of the sub-service at Line 35. The (blocking) invocation will return only when the execution of the scope is terminated.

```

1  outputPort cinema {
2      Location: "socket://localhost:10500/"
3      Protocol: sodep
4      RequestResponse: op0(OpType)(undefined)
5  }
6
7  outputPort alice {
8      Location: "socket://localhost:10501/"
9      Protocol: sodep
10     RequestResponse: op0(OpType)(undefined)
11 }
12
13 main
14 {
15     var12.role = "f241629d94f2";
16     start@MH(var12)();
17     start_bob@MH(var12)();
18     end = false;
19     var0.value = end;
20     var0 = "end";
21     set@State(var0)();
22     get@State("end")(end);
23     var7 = !end;
24     var8.role = "709bdb0d6192";
25     var8.content = var7;
26     op0@cinema(var8)();
27     var9.role = "709bdb0d6192";
28     var9.content = var7;
29     op0@alice(var9)();
30     while ( var7 ){
31         run@ActivityManager("0a8a09022eee")();
32         var2.role = "5355f138ae39";
33         get_op1@MH(var2)(var2);
34         if (var2.content) {
35             run@ActivityManager("c7453ecb0abf")()
36         };
37         get@State("end")(end);
38         var6 = !end;
39         if (var6) {
40             showInputDialog@SwingUI("Alice refused." +
41                 "Try to propose another day?[y/n]" )(_r);
42             var3.value = _r;
43             var3 = "_r";
44             set@State(var3)();
45             get@State("_r")(_r);
46             var5 = _r != "y";
47             if (var5) {
48                 end = true;
49                 var4.value = end;
50                 var4 = "end";
51                 set@State(var4)()
52             }
53         };
54         var10.role = "cinema";
55         get_op0@MH(var10)();
56         var11.role = "alice";
57         get_op0@MH(var11)();
58         get@State("end")(end);
59         var7 = !end;
60         var8.role = "709bdb0d6192";
61         var8.content = var7;
62         op0@cinema(var8)();
63         var9.role = "709bdb0d6192";
64         var9.content = var7;
65         op0@alice(var9)()
66     }
67 }

```

Listing 1.5: Bob Jolie code.

Fig. 6 shows the execution of the choreography (with one service of the choreography per console). Consoles show also some debugging information to clarify the flow of execution. The application is waiting for receiving the data that will be inserted in variable `free_date` by Bob (Line 12, Listing 1.1).

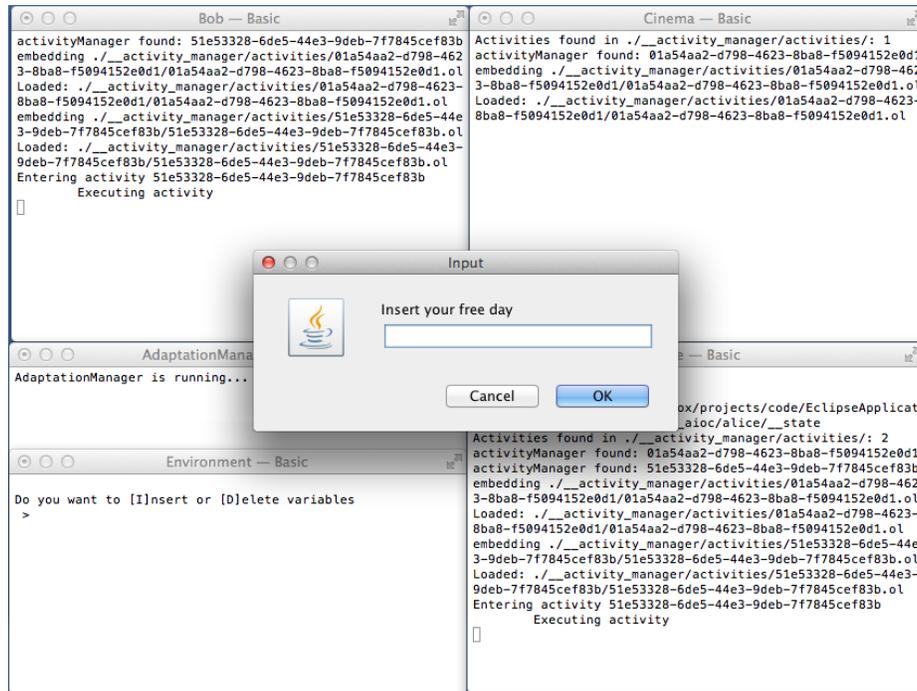


Fig. 6: Start of the AIOCJ application.

In this execution, Alice is free the day Bob proposes, but does not agree with going to the cinema (Fig. 7). Bob is notified of the refusal and asked if he wants to retry, proposing another day.

A.4 Writing and compiling adaptation rules

While the AIOCJ application is running, one can write, compile, and publish new rules.

The rule in Listing 1.6 adapts the scope named `"matching day"` at Lines 11-15 of Listing 1.1. It changes the protocol used by Bob to ask Alice for the day of the appointment. Bob now first checks, using the external function `hasFreeWeek`, if he is free the whole next week. If so, he asks Alice to pick a day, otherwise he follows the normal protocol, suggesting to Alice a day for an appointment when he is free.

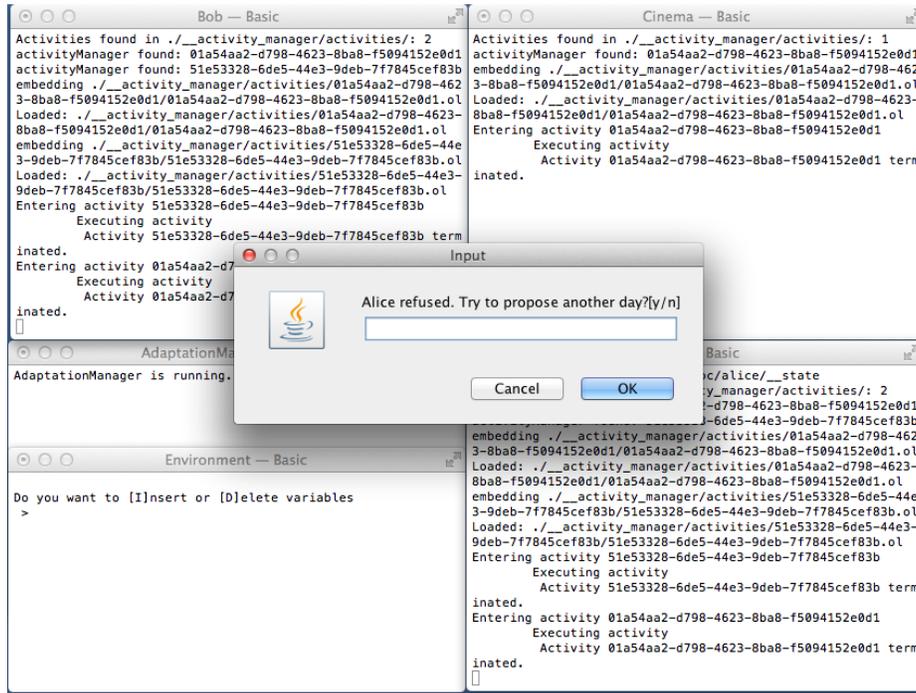


Fig. 7: First iteration, refusal.

```

1 rule {
2   include isFreeDay, hasFreeWeek from "socket://localhost:8000"
3   on { N.scope_name == "matching day" }
4   do {
5     not_busy@bob = hasFreeWeek();
6     if ( not_busy )@bob {
7       free_day@alice = getInput( "Bob has a free week,
8         insert your free day or no to refuse" );
9       if (free_day == "no")@alice {
10        is_free@alice = false;
11      } else {
12        is_free@alice = true;
13        response: alice( free_day ) -> bob( free_day ) }
14    } else {
15      free_day@bob = getInput( "Insert your free day" );
16      proposal: bob( free_day ) -> alice( bob_free_day );
17      is_free@alice = isFreeDay( bob_free_day ) }}}

```

Listing 1.6: Matching day adaptation rule.

Rules are compiled following the same procedure of AIOC applications. The compilation creates a new folder “`app_rules`” that contains an adaptation server and the compiled rules.

A.5 Adapting an AIOcJ application

Inside the folder “epp_rules” the script “rules_launcher.sh” starts the adaptation server. Once started, the adaptation server registers itself to the adaptation manager. If Bob decides to proceed with a new iteration of the choreography, the application adapts. The rule in Listing 1.6 applies. Fig. 8 shows the additional console where the adaptation server is running (bottom left corner). The dialog box confirms that the rule in Listing 1.6 has been applied (it notifies to Alice that Bob has a free week).

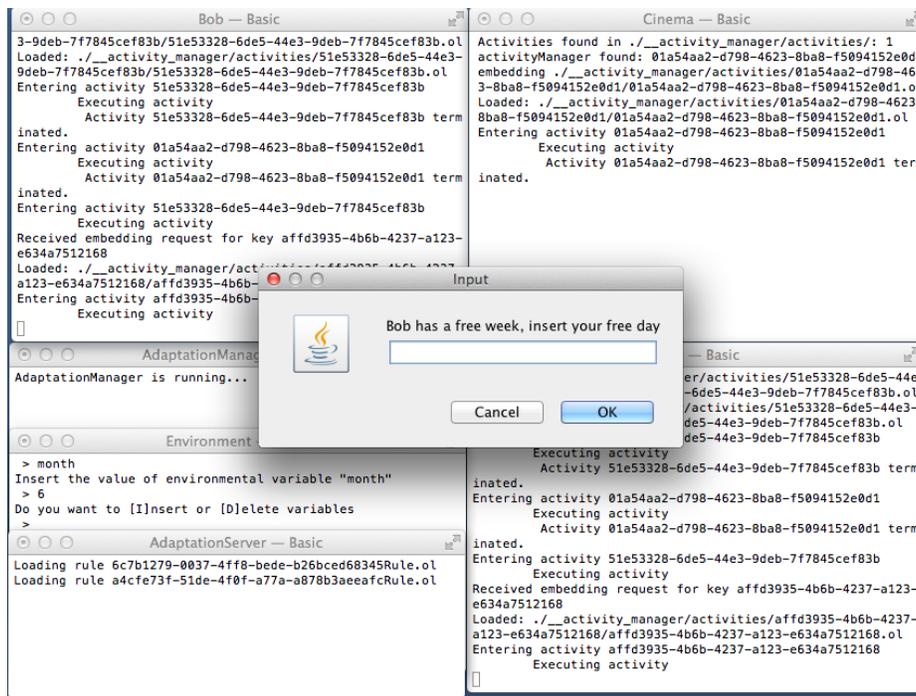


Fig. 8: Adaptation using the “matching day” rule.

In order to apply also the rule in Listing 1.2, we set the variable month in the environment to 6 and proceed with the choreography. Fig. 9 shows the result of the adaptation in which Bob asks whether Alice wants to go out for a picnic.

The code of the services that implement functions isFreeDay and getTicket is given in [27], where this example is hosted, including the AIOC program, the rules, and the additional services.

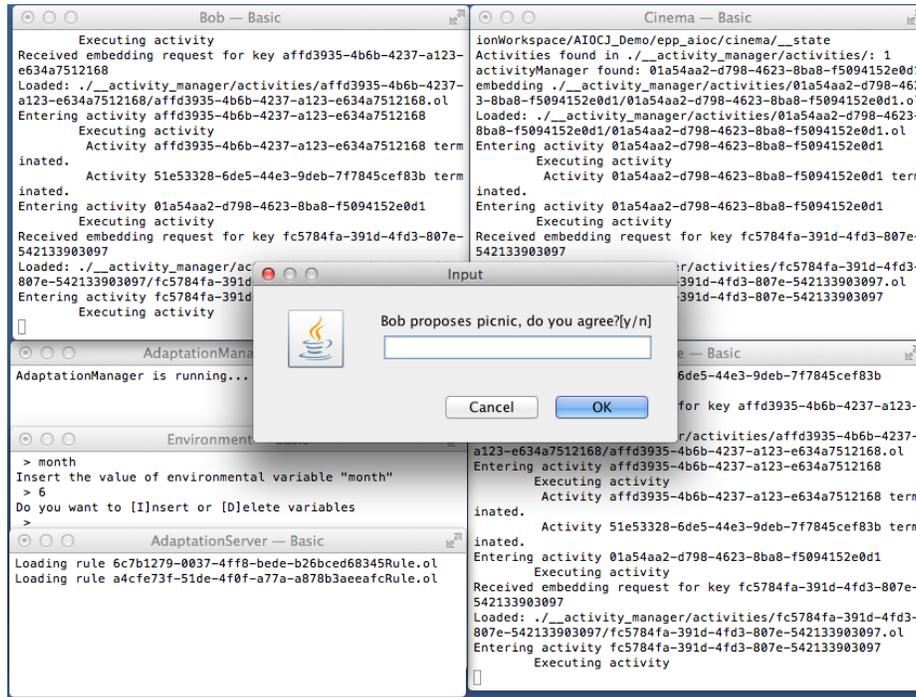


Fig. 9: Second iteration, adapted execution.

B Code used for validation

B.1 Pipe and fork-join code

In this section, we provide the code of the AIOCJ programs used in Section 4. For simplicity, we show the programs fixing the number of tasks $n = 5$, and showing just two adaptation rules per scenario.

Pipe scenario. In the pipe scenario each task computes the increment function, passing the output of its computation as input to the next task.

```

1 include startTimer, endTimer from "socket://localhost:8000"
2 preamble{ starter: a }
3
4 aioc {
5   { x@a = 0 | x@b = 0 };
6   _r@a = startTimer( n );
7   x@a = 1 + x; pass: a( x ) -> b( x );
8   x@b = 1 + x; pass: b( x ) -> a( x );
9   x@a = 1 + x; pass: a( x ) -> b( x );
10  x@b = 1 + x; pass: b( x ) -> a( x );
11  x@a = 1 + x; pass: a( x ) -> b( x );
12  _r@a = stopTimer( n )
13 }

```

Listing 1.7: Pipe scenario without scopes.

Listing 1.7 shows the 5 tasks that increment variable `x`. After each increment the variable is sent from role `a` to role `b` or vice versa. Lines 6 and 12 are calls to the `Timer` service (Line 1), used to log the times of execution.

```
1 include startTimer, endTimer from "socket://localhost:8000"
2 preamble{ starter: a }
3
4 aioc {
5   { x@a = 0 | x@b = 0 };
6   _r@a = startTimer( n );
7   scope @a {x@a = 1 + x; pass: a( x ) -> b( x )} prop {N.x = 1};
8   scope @b {x@b = 1 + x; pass: b( x ) -> a( x )} prop {N.x = 2};
9   scope @a {x@a = 1 + x; pass: a( x ) -> b( x )} prop {N.x = 3};
10  scope @b {x@b = 1 + x; pass: b( x ) -> a( x )} prop {N.x = 4};
11  scope @a {x@a = 1 + x; pass: a( x ) -> b( x )} prop {N.x = 5};
12  _r@a = stopTimer( n )
13 }
```

Listing 1.8: Pipe scenario with scopes.

Listing 1.8 shows the 5 scopes enclosing the tasks of Listing 1.7. Notably, each scope has a unique property `N.x` assigned incrementally. This property is used to identify uniquely each scope.

Rules of the pipe scenario. Listing 1.9 shows two rules used in the pipe scenario. The adapted behaviour increments the variable `x` by 2. In particular, the rule at Lines 1-4 applies to the scope at Line 7 in Listing 1.8, while the subsequent rule (Lines 5-8) applies to the scope at Line 8.

```
1 rule {
2   on { N.x == 1 }
3   do { x@a = 2 + x; pass: a( x ) -> b( x ) }
4 }
5 rule {
6   on { N.x == 2 }
7   do { x@b = 2 + x; pass: b( x ) -> a( x ) }
8 }
```

Listing 1.9: Rules for pipe scenario.

Fork-join scenario. In the fork-join scenario tasks run in parallel. Each task takes one character of a message and replaces it with the next character in the alphabet. The program makes use of an external service to retrieve the character of the message, get the next letter, and store the new character in the message.

```

1 include start, end from "socket://localhost:8000"
2 include getNthChar, getNext, setNthChar from "socket://localhost:8001"
3
4 preamble{ starter: a }
5
6 aioc {
7   _rea = start( n );
8   {
9     { 10@a = getNthChar( 0 ); 10@a = getNext( 10 ); _rea = setNthChar( 0, 10 ) }
10    | { 11@b = getNthChar( 1 ); 11@b = getNext( 11 ); _rea = setNthChar( 1, 11 ) }
11    | { 12@a = getNthChar( 2 ); 12@a = getNext( 12 ); _rea = setNthChar( 2, 12 ) }
12    | { 13@b = getNthChar( 3 ); 13@b = getNext( 13 ); _rea = setNthChar( 3, 13 ) }
13    | { 14@a = getNthChar( 4 ); 14@a = getNext( 14 ); _rea = setNthChar( 4, 14 ) }
14    };
15   _rea = end( n )
16 }

```

Listing 1.10: Fork-join scenario without scopes.

Listing 1.10 shows the 5 tasks of the fork-join scenario. Each task uses a dedicated variable (10, 11, ..., 1n) to ensure thread-safety of the computation.

```

1 include start, end from "socket://localhost:8000"
2 include getNthChar, getNext, setNthChar from "socket://localhost:8001"
3
4 preamble{ starter: a }
5
6 aioc {
7   _rea = startTimer( n ); {
8     scope @a {
9       10@a = getNthChar( 0 ); 10@a = getNext( 10 ); _rea = setNthChar( 0, 10 ) }
10      prop { N.char = 0 } |
11      scope @b {
12        11@b = getNthChar( 1 ); 11@b = getNext( 11 ); _rea = setNthChar( 1, 11 ) }
13      prop { N.char = 1 } |
14      scope @a {
15        12@a = getNthChar( 2 ); 12@a = getNext( 12 ); _rea = setNthChar( 2, 12 ) }
16      prop { N.char = 2 } |
17      scope @b {
18        13@b = getNthChar( 3 ); 13@b = getNext( 13 ); _rea = setNthChar( 3, 13 ) }
19      prop { N.char = 3 } |
20      scope @a {
21        14@a = getNthChar( 4 ); 14@a = getNext( 14 ); _rea = setNthChar( 4, 14 ) }
22      prop { N.char = 4 }
23    };
24   _rea = stopTimer( n )
25 }

```

Listing 1.11: Fork-join scenario with scopes.

Likewise the pipe scenario, the scopes in Listing 1.11 enclose the tasks of Listing 1.10. Also in this scenario each scope has a unique property `N.char` that identifies it uniquely.

Rules of the fork-join scenario. The adapted behaviour introduces the new function `GetDoubleNext` that gets a character and returns the character two positions after it in the alphabet. Listing 1.12 shows two rules used in the fork scenario. The rule at Lines 1-6 applies to the scope at Lines 8-10 of Listing 1.11, while the rule at Lines 7-12 applies to the scope at Lines 11-13.

```

1 rule {
2   include getNthChar, getDoubleNext, setNthChar from "socket://localhost:8001"
3   on { N.char == 0 }
4   do { l0@b = getNthChar( 0 ); l0@b = getDoubleNext( 10 );
5     _r@b = setNthChar( 0, 10 ) }
6 }
7 rule {
8   include getNthChar, getDoubleNext, setNthChar from "socket://localhost:8001"
9   on { N.char == 1 }
10  do { l1@b = getNthChar( 1 ); l1@b = getDoubleNext( 11 );
11    _r@b = setNthChar( 1, 11 ) }
12 }

```

Listing 1.12: Rules of pipe scenario.

B.2 AIOCJ programs used for benchmarking primitives

```

1 include start, end from "socket://localhost:8000"
2 preamble { starter: a }
3 aioc {
4   _r@a = start( "assignment" );
5   x@a = 1;
6   _r@a = end( "assignment" );
7
8   _r@a = start( "interaction" );
9   pass: a( x ) -> b( x );
10  _r@a = end( "interaction" );
11
12  _r@a = start( "if statement" );
13  if ( x == 1 )@a { skip };
14  _r@a = end( "if statement" )
15 }

```

Listing 1.13: Code for benchmarking assignment, interaction, and the if statement.

```

1 include start, end
2   from "socket://localhost:8000"
3
4 preamble { starter: a }
5
6 aioc {
7   _r@a = start( "scope" );
8   scope @a{ skip }
9   prop { N.scope_name = "applicable" };
10  _r@a = end( "scope" )
11 }

```

Listing 1.14: Code for benchmarking the scope primitive.

Listings 1.13 and 1.14 show the code used to measure the performances of primitives of AIOCJ. We tested the performances of the scope (Listing 1.14) under 7 contexts:

1. without adaptation servers;
2. with 1 adaptation server, but without rules;
3. with 1 adaptation server and 1 matching rule;
4. with 1 adaptation server, 50 rules but none matching;
5. with 1 adaptation server, 50 rules, and 1 matching;
6. with 1 adaptation server, 100 rules but none matching;

7. with 1 adaptation server, 100 rules, and 1 matching;

```
1 rule {
2   on { N.scope_name == "applicable" }
3   do { skip }
4 }
5 rule {
6   on { N.scope_name == "non-applicable" }
7   do { skip }
8 }
```

Listing 1.15: Rules used for benchmarking scopes.

Listing 1.15 shows the two kinds of rules used for testing performances of Listing 1.14. The first one (Lines 1-4) is the matching rule used in contexts 3, 5, and 7. The second one (Lines 5-8) is a non matching rule present, in different copies, in contexts 4, 5, 6, and 7.