

Multiparty Languages: The Choreographic and Multitier Cases

Saverio Giallorenzo   

Department of Computer Science and Engineering, Università di Bologna, Italy and INRIA, France

Fabrizio Montesi   

Department of Mathematics and Computer Science, University of Southern Denmark, Denmark

Marco Peressotti   

Department of Mathematics and Computer Science, University of Southern Denmark, Denmark

David Richter 

Technical University of Darmstadt, Germany

Guido Salvaneschi  

University of St.Gallen, Switzerland

Pascal Weisenburger  

Technical University of Darmstadt, Germany

Abstract

Choreographic languages aim to express multiparty communication protocols, by providing primitives that make interaction manifest. Multitier languages enable programming computation that spans across several tiers of a distributed system, by supporting primitives that allow computation to change the location of execution. Rooted into different theoretical underpinnings—respectively process calculi and lambda calculus—the two paradigms have been investigated independently by different research communities with little or no contact. As a result, the link between the two paradigms has remained hidden for long.

In this paper, we show that choreographic languages and multitier languages are surprisingly similar. We substantiate our claim by isolating the core abstractions that differentiate the two approaches and by providing algorithms that translate one into the other in a straightforward way. We believe that this work paves the way for joint research and cross-fertilisation among the two communities.

2012 ACM Subject Classification Computing methodologies → Distributed programming languages; Theory of computation → Distributed computing models; Software and its engineering → Multiparadigm languages; Software and its engineering → Concurrent programming languages; Software and its engineering → Distributed programming languages

Keywords and phrases Distributed Programming, Choreography Programming, Multitier Programming

Digital Object Identifier 10.4230/LIPIcs.ECOOP.2021.23

Category Pearl

Funding *Fabrizio Montesi*: Villum Fonden, grant no. 29518, and Independent Research Fund Denmark, grant no. 0135-00219.

Marco Peressotti: Villum Fonden, grant no. 29518, and Independent Research Fund Denmark, grant no. 0135-00219.

1 Introduction

Programming concurrent and distributed systems is notoriously hard. Among other issues, it requires dealing with coordination and predicting how multiple participants will interact at



© Saverio Giallorenzo, Fabrizio Montesi, Marco Peressotti, David Richter, Guida Salvaneschi, and Pascal Weisenburger;

licensed under Creative Commons License CC-BY 4.0

35th European Conference on Object-Oriented Programming (ECOOP 2021).

Editors: John Q. Open and Joan R. Access; Article No. 23; pp. 23:1–23:27

Leibniz International Proceedings in Informatics



Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

43 runtime, for which programmers do not receive adequate help from mainstream programming
44 abstractions and technology [25, 21, 32].

45 The quest for finding elegant languages and methodologies that can help with concurrent
46 and distributed programming has been a major focus of the research community for decades,
47 including the seminal actor model and calculus of communicating systems [17, 27]. In this
48 work, we are interested in two kinds of languages that have been recently gaining attention:
49 *choreographic languages* [28, 2] and *multitier languages* [40]. Choreographic languages
50 are designed to express multiparty communication protocols, by providing primitives that
51 make interaction manifest. On the other hand, multitier languages allow for programming
52 computation that spans across several tiers of a distributed system, by providing primitives
53 that allow computation to change location of execution.

54 Both choreographic and multitier languages aim at making concurrent and distributed
55 programming more effective, and have inspired several research and industrial language
56 designs. However, choreographic and multitier languages stem from different ideas; they
57 adopt different terminologies; they look different; they have evolved different features; and
58 they have found different applications in practice. Perhaps because the design principles of
59 choreographic and multitier languages come from different angles, the two communities have
60 prolifically evolved independently. However, as a consequence, the commonalities and actual
61 differences between the two research lines remain unclear, which impedes cross-fertilisation.

62 In this paper, we offer a new perspective on the relationship between choreographic and
63 multitier languages. We show that, despite their different starting points and evolutions,
64 they share a strong core idea that classifies them both as what we call *multiparty languages*—
65 languages that describe the behaviour of multiple participants. Leveraging this commonality,
66 it is possible to derive choreographic programs from multitier programs, and vice versa. Our
67 aim is to provide a way for each community to access the other, encouraging cross-fertilisation.

68 We outline our investigation and contributions:

- 69 ■ In Section 2, we give an overview of the essential features of choreographic and multitier
70 languages. We recap the history of the two approaches and identify their key differences,
71 which lie in perspective (objective vs subjective) and in the modelling of communications
72 (manifest vs non-manifest). We also pinpoint the commonality that classifies choreographic
73 and multitier languages as multiparty.
 - 74 ■ In Section 3, we present an example use case for both choreographic and multitier
75 programming, which introduces the concrete choreographic and multitier programming
76 languages that we will use in the rest of our development: Choral [16] and ScalaLoc [38].
 - 77 ■ In Section 4, we introduce Mini Choral and Mini ScalaLoc, two representative but
78 minimal languages for choreographic and multitier programming, respectively. Mini
79 Choral and Mini ScalaLoc dispense with the features that are not essential parts of
80 their respective paradigms, which allows us to study how the essential differences can be
81 bridged in the next section.
 - 82 ■ In Section 5, we define algorithms for translating programs in Mini Choral to programs in
83 Mini ScalaLoc, and vice versa. The translations deal with the changes in perspective and
84 manifestation of communications between the two paradigms. For example, translating
85 a multitier program into a choreographic one requires synthesising a communication
86 protocol that enacts the necessary communications among participants.
- 87 Our translations are not just of inspiration to see the connection between the two
88 paradigms (which we leverage in the next section), but also open a window towards the
89 future sharing of theoretical and practical results. An example for each direction: by
90 translating a multitier program into a choreographic one and then using a choreographic

91 compiler to generate executable code, we can know statically the pattern of communica-
 92 tions that will be enacted by the executable code (this property is called “Choreography
 93 Compliance” [16] or “EndPoint Projection Theorem” [3]); by translating a choreographic
 94 program into a multitier one and then using a multitier compiler to generate executable
 95 code, we can reuse all the machinery developed by the multitier community to generate
 96 code for different technologies (e.g., the code generated for one participant is in JavaScript
 97 for a web browser while the code for another might be code runnable on the Java Virtual
 98 Machine for a server).

99 ■ Our study shows that, while choreographic and multitier programming languages are
 100 different enough to be independently useful, they are also near enough to benefit from
 101 cross-fertilisation. In Section 6, we report on important features that have been developed
 102 separately in the choreographic and multitier research lines. We find that important
 103 features for the development of concurrent and distributed systems have been developed
 104 for one paradigm but not the other. Inspired by our newfound connection, we discuss
 105 how these features could be ported over to the other paradigm in the future, setting up
 106 future work enabled by our view.

107 2 Background: Choreographic and Multitier Programming Languages

108 In this section, we give some background on choreographic and multitier languages, and
 109 discuss their differences and similarities.

110 2.1 Choreographic Languages

111 Choreographic languages are inspired by the famous “Alice and Bob” notation, or security
 112 protocol notation [30]. The idea is to define how the different participants of a system should
 113 communicate (or interact)—which later inspired also message sequence charts and sequence
 114 diagrams [20]. Textual and graphical choreographic languages have already been adopted in
 115 industry as specification languages in different settings ranging from business processes, e.g.,
 116 the choreographic language in OMG’s Business Process Model and Notation, to web services,
 117 e.g., W3C’s Web Services Choreography Description Language [31, 37].

118 The essence of a choreographic language is the capability of expressing explicitly data flows
 119 from a participant to another through communication, and of composing such communications
 120 into larger structures. In other words, choreographies make interaction and the structure of
 121 interaction protocols *manifest*. A communication from a participant, **Alice**, to another, **Bob**,
 122 is written as follows.

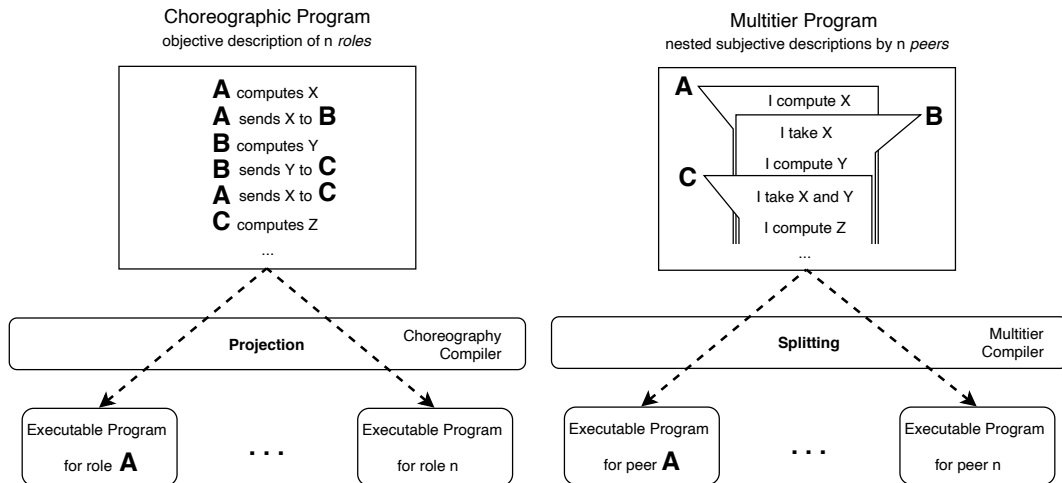
```
123 Alice.userId -> Bob.x : ch
```

126 The statement above reads: **Alice** sends its `userId` (a local variable storing a user identifier)
 127 to **Bob**, which stores it in its local variable `x`, and the communication takes place through the
 128 channel `ch`.

129 Communication statements can be composed in larger and more sophisticated protocols,
 130 for example using the sequential operator “;”. In the following protocol snippet: after
 131 interacting with **Alice**, **Bob** forwards to **Charlie** the user identifier that it received through a
 132 separate channel `ch2`.

■ **Listing 1** A simple choreography with three participants.

```
133 Alice.userId -> Bob.x : ch;  
134 Bob.x -> Charlie.y : ch2;
```



■ **Figure 1** Choreographic Programming.

■ **Figure 2** Multitier Programming.

137 In the paradigm of *choreographic programming* [28], choreographic languages are full-
 138 fledged programming languages: developers write the implementation of an entire multiparty
 139 system as a choreography, and then a compiler automatically generates an executable
 140 program for each participant. This process is depicted in Figure 1. Choreographies resemble
 141 play scripts, written from an external point of view, describing the interactions among all
 142 participants. We call this view *objective*. Participants, like **Alice** and **Bob**, are typically
 143 referred to as *roles* in choreographies, and the procedure that generates the executable
 144 program for each role is called projection (or endpoint projection) [4, 11].

145 The code in Listing 1 is valid code in the Chor language, the first implementation of
 146 choreographic programming [28, 4]. Chor targets microservices: given that code (with
 147 appropriate boilerplate), Chor would generate executable programs of microservices that
 148 implement **Alice**, **Bob**, and **Charlie**. Choreographic programming has been applied to other
 149 settings, e.g., information flow [22], parallel algorithms [10], cyber-physical systems [24, 23],
 150 runtime adaptation [11], and integration processes [15].

151 2.2 Multitier Languages

152 Multitier languages are inspired by one of the ideas proposed with ambient calculi [5]. In this
 153 kind of process calculi, terms express the place (the “ambient”) at which computation occurs.
 154 Computations that take place at different locations can be nested, which enables describing
 155 multiparty systems. It was later shown that the idea can be combined with well-known
 156 abstractions, by developing a variation of λ -calculus with locations called Lambda 5 [29].
 157 This solution prompted the development of *multitier languages* [36, 8, 40], which extend
 158 existing programming languages with locations. The term multitier comes from the fact that
 159 these languages were mostly developed for web programming, where tiers is used to refer to
 160 the typical participants of a web system (e.g., client, backend server, and database).

161 The crux of a multitier language is the capability of hopping from the point of view of
 162 a participant to that of another—the multitier language by Serrano et al. is aptly called
 163 “Hop” [36]. When hopping from a participant to another, it is possible to move data
 164 from the participant that we are leaving to the participant that we are going to—enabling
 165 communication. As an example, consider a remote procedure call from a client to server. In a

166 recent incarnation of multitier programming that builds on the Scala language, ScalaLoci [38],
 167 this can be written as follows.¹

```

168 def rpc(input: String): String on Client = on[Client] {
169   val result =
170     on[Server].run.capture(input) {
171       expensiveFunction(input)
172     }.asLocal
173   return result
174 }
175

```

177 Participants are referred to as peer types in ScalaLoci. The method `rpc` above is defined
 178 as a block of code that starts at the client peer (`on[Client]`). The client stores the result
 179 of some computation in its local variable `result`, but this computation is performed at the
 180 server. This result is achieved by “moving” to the server with the instruction `on[Server]`. The
 181 invocation of method `run`, right afterwards, models some computation, and `capture(input)`
 182 means that we want to move the content of the local variable `input` from the peer that we are
 183 leaving (the client) to the one that we are going to (the server). How this move is achieved
 184 is left to the implementation (ScalaLoci generates a communication strategy automatically).
 185 The server then runs an expensive function on the input, and the execution goes back to the
 186 client—the code block at the server ends. The invocation of `asLocal` ensures that the return
 187 value of the code at the server is moved to the location of the enclosing scope (the client).
 188 We finally return the result at the client.

189 Like choreographic programming languages, multitier languages come with a compiler
 190 that turns the multiparty view of the system into executable programs. This process is
 191 depicted in Figure 2. Given a multitier program, a multitier compiler generates an executable
 192 program for each peer type (in the case of Section 2.2, these would be client and server). The
 193 procedure for generating code is called splitting. The nested “dialogues” of peers inside the
 194 multitier program depict that a multitier program has many viewpoints, switching regularly
 195 from the point of view of a peer to that of another. Nevertheless, code is written with the
 196 viewpoint of the peer we are currently in. For this reason, we say that multitier programs
 197 adopt a nested *subjective* view.

198 2.3 Towards Linking Choreographic to Multitier Languages

199 The two communities of choreographic and multitier languages have prolifically evolved
 200 independently [2, 40]. They adopted different design principles, and they have found different
 201 practical applications—most notably service-oriented computing for choreographies and
 202 web development for multitier programming. As a result, they have also developed several
 203 features independently (we discuss some of the most important ones in Section 6). In
 204 addition, the two communities have been facing different challenges. For example, multitier
 205 programming languages historically tackle the problem of “impedance mismatch”: the
 206 necessity of handling data conversions and heterogeneous execution engines in the web
 207 (the Google Web Toolkit is a multitier framework that contributes to this research area).
 208 Instead, choreographic programming mainly aimed at achieving “choreography compliance”:
 209 providing the guarantee that distributed systems communicate as expected and with desirable
 210 properties (like liveness).

211 Yet, the two paradigms are clearly linked. We drew Figure 1 and Figure 2 with the
 212 intention of highlighting such connection. Indeed, despite differences in both terminologies

¹ For simplicity of presentation, we omit library calls that would be necessary to deal with asynchrony.

23:6 Multiparty Languages: The Choreographic and Multitier Cases

and methods, the strategies of choreographic and multitier programming languages share a similarity: both define the behaviour of a multiparty system in a single compilation unit, and then offer ways to synthesise executable implementations for the participants. We thus identify both kinds of languages as instances of the larger class of *multiparty languages*—leaving the class open to future additions. We see value in both techniques for multiparty programming. In choreographies protocols are manifest, which makes them easy to understand. Multitier programs give access to multiparty programming with a developing experience that resembles standard “local programming” by leveraging scoping.

Despite both choreographic and multitier languages sharing the multiparty approach, they remain pretty diverse in terms of theoretical background. The theory of choreographic language typically stands on process calculi, whereas multitier models build on λ -calculus [18, 4, 19, 11, 40]. This is likely an important reason why the link between choreographic and multitier languages has been overlooked for long. Very recently, however, it has been shown that object-oriented languages can be extended to capture choreographies, by generalising the notion of data type to data types located at *multiple roles* [16]. In the resulting language, called Choral, a choreography among a few roles can be expressed as an object. For example, we can write the choreography in Listing 1 in Choral as follows:

```
230
2311 class Example@(Alice, Bob, Charlie) { // the three roles of the protocol
2322   DiDataChannel@(Alice,Bob)<Serializable> ch; // channel from Alice to Bob
2333   DiDataChannel@(Bob,Charlie)<Serializable> ch2; // channel from Bob to Charlie
2344
2355   /* constructor omitted */
2366
2377   public UserID@Charlie run(UserID@Alice userId) { // the protocol
2388     UserID@Bob x = ch.<UserID>com(userId); // Alice.userId -> Bob.x : ch
2399     return ch2.<UserID>com(x); // Bob.x -> Charlie.y : ch2
2400   }
2411 }
```

Briefly—as we give a more detailed description of Choral programs in Section 3.2—the `Example` class declares three roles (`Alice`, `Bob`, and `Charlie`) and two directed channels (`ch` from `Alice` to `Bob` and `ch2` from `Bob` to `Charlie`). These correspond to the roles and channels assumed in Listing 1. The protocol described in Listing 1 is implemented by method `run` that takes an instance of `UserID` located at `Alice` and returns one located at `Charlie` passing through `Bob`. Communication happens by invoking method `com` of the two channels.


Choral helps in leveling the playfield with multitier programming. Indeed, we now have an object-oriented incarnation of choreographic programming that we can use to compare to object-oriented multitier languages, here represented by `ScalaLoci`. In the next sections, we leverage this common ground and take Choral and `ScalaLoci` as representative languages for their respective paradigms.

3 Overview of Choral and ScalaLoci

In this section, we give an overview of the representative languages for choreographic and multitier programming that we have chosen, Choral and `ScalaLoci`, by using them to deal with a simple yet comprehensive example of a context-aware protocol for e-mail fetching.

3.1 A Context-Aware Email-Fetching Protocol

Before delving into the details of the two implementations, we discuss briefly the protocol that we want to program. A depiction as a sequence diagram is given in Figure 3. The protocol defines an interaction between an Email Client and an Email Server. Specifically,



■ **Figure 3** Sequence diagram for context-aware e-mail fetching.

262 the Client sends its identification token—here simplified as User ID—and the timestamp of
 263 the last e-mail checkout to the Server. The Server returns the list of e-mails received after
 264 the timestamp to the Client. After the above interaction, the Client and the Server enter
 265 an optional block. The optional block is executed depending on the context of the client,
 266 namely, if the connection from the Client to the Server is flat-rate, i.e., if the connection fee
 267 paid by the Client is independent from its usage. If that is the case, the Client sends the
 268 Server the list of e-mail IDs retrieved in the previous interaction to fetch their attachments.
 269 The Server concludes the optional part of the protocol by sending to the Client the requested
 270 attachments.

271 3.2 A Choreographic Programming Implementation with Choral

272 In Listing 2, we use Choral to implement the protocol from Figure 3. The example illustrates
 273 the main concepts of the choreographic programming approach and how Choral captures
 274 them in the object-oriented setting.

275 In Choral, objects have types of the form $\tau@(\mathbf{R1}, \dots, \mathbf{Rn})$, where τ is the interface of the
 276 object (as usual), and $\mathbf{R1}, \dots, \mathbf{Rn}$ are the roles that collaboratively implement the object.
 277 As we see below, Choral supports two notations for object ownership: the standard form
 278 $@(\mathbf{A}, \dots, \mathbf{Z})$ and the contracted form $@\mathbf{A}$, for objects that belong to one role (shortcut for
 279 $@(\mathbf{A})$). Incorporating roles in data types makes distribution manifest at the type level.

280 In Listing 2, at Line 3, we define a class `EmailSystem` implemented by two roles: the `Client`
 281 and the `Server`. The method `updateEmails` (Line 8) implements the actual protocol from
 282 Figure 3. Lines 4–6 declare class-level `private` objects, i.e., accessible from the `updateEmails`
 283 method and other (omitted) ones within the class. Specifically, at Line 4, we have the
 284 `MailServerDB` located at the `Server`. At Line 5, we find the complementary `MailDB` of the
 285 `Client`. At Line 6, we define the object used to transfer data between the two roles: a
 286 `SymChannel`—standing for *symmetric channel*—shared between the two roles and able to
 287 transmit `Serializable` objects. We omit the initialisation of the abovementioned objects.

288 Considering the description of the implementation of the e-mail fetching protocol, we look
 289 at the `updateEmails` method (Line 8). The method does not return a value (`void`) and takes

23:8 Multiparty Languages: The Choreographic and Multitier Cases

290 as input the `UserId`—which simplifies the user authentication procedure here, for brevity—to
291 identify the user of the `Client` at the `Server`.

292 In the body, at Line 9, we pass the `UserId` to the `Server`. We do this by invoking the
293 method `com` of the `ch SymChannel` giving to it as argument the `userId`. This is done by the
294 expression `userId >> ch::com` which uses the Choral chaining operator `>>` and that corresponds
295 to the expanded expression `ch.com(userId)`.² The method `com` of the `SymChannel` transfers the
296 value of the sender given as input into an equivalent representation of the value at the receiver.
297 In this case, the sender is the `Client` (where the `UserId` object lives) and the receiver is the
298 `Server`, which stores the result of the communication into variable `id` which is an object of
299 type `UserId` at its location—i.e., `UserId@Server`.

300 The transfer of the `Timestamp` from the `Client` to the `Server` is similar (Line 10): we retrieve
301 the object from the `clientDB`—invoking method `lastCheckout`—and we transfer it to the `Server`
302 through the `SymChannel`. Then, to fetch the e-mails, the `Client` receives a transmission from
303 the `Server`. The `Server` interrogates its local database (`serverDB`) by extracting all e-mails
304 belonging to the `id` of the `Client` and received since its last checkout (indicated by the
305 `timestamp`) and sends them to the `Client` via their shared `SymChannel`. At Line 12, the `Client`
306 uses the received list of emails to update its local database (`clientDB`).

307 Lines 13–20 implement the optional part of the protocol from Figure 3. First, the `Client`
308 checks whether it is using a flat-rate connection—this is done through the static library
309 `ClientLib` and its method `isOnFlatRate`.

310 The `if-else` block at Lines 13–20 allows us to explain the concept of knowledge of choice
311 (a hallmark element of choreographic programming) and how Choral implements it. Briefly,
312 the concept of knowledge of choice indicates a fork in the flow of a program among alternative
313 behaviours, where the concerned roles should coordinate to ensure that they agree on which
314 behaviour they should enact. In Choral, we adopt a standard choreographic solution to this
315 problem [11] by defining a “selection” primitive to communicate constants drawn from a
316 dedicated set of “labels”, so that the compiler has enough information to build code that
317 can react to choices made by other roles. Concretely, to define selections, Choral uses a
318 method-level annotation `@SelectionMethod`³, which developers can apply only to methods
319 that can transmit instances of enumerated types between roles (the compiler checks for this
320 condition). Conveniently, the `SymChannel` used in the example also supports selections via its
321 `select` methods. In Listing 2, we find the implementation of the knowledge of choice of the
322 conditional at Line 14 (where the `Client` “decides” to fetch the attachments) and at Line 20
323 (which skips the retrieval). In the example, we implement the choice by defining the `Choice`
324 `enum` class at Line 1—note that we use the identifier `Role` for the single role that owns the
325 `Choice` object in its declaration, instantiated at the `Client` at Lines 14 and 20.

326 If the `Client` uses a flat-rate connection, the chained statement at Lines 15–17 execute:
327 first (Line 15) the `Client` sends to the `Server` the IDs of the e-mails (retrieved through
328 `extractIds(emails)`) whose attachments it wants to retrieve, then (Line 16) the `Server` uses
329 the received `ids` to extract from its database (`serverDB`) the attachments and it send them
330 back to the `Client`, and finally (Line 17) the `Client` uses the received attachments to update
331 its local database.

² To make Choral programs closer to standard choreographic notation, where data flows from left to right, Choral borrows the forward chaining operator `>>` from F#: `exp >> obj::method` is syntactic sugar for `obj.method(exp)`.

³ Choral preserves the standard `@`-notation for annotations from Java, which is contextually separated from `@(R1, ..., Rn)` parameters in Choral programs.

■ **Listing 2** Choral implementation for the context-aware e-mail fetching example.

```

1  enum Choice@Role { THEN, ELSE }
2
3  class EmailSystem@(Client, Server) {
4    private MailServerDB@Server serverDB = ...;
5    private MailDB@Client clientDB = ...;
6    private SymChannel@(Client, Server)<Serializable> ch = ...;
7
8    void updateEmails(UserId@Client userId) {
9      UserId@Server id = userId >> ch::com;
10     Timestamp@Server timestamp = clientDB.lastCheckout() >> ch::com;
11     List@Client<Email> emails = serverDB.since(id, timestamp) >> ch::com;
12     clientDB.update(emails);
13     if (ClientLib@Client.isOnFlatRate()) {
14       Choice@Client.THEN >> ch::select;
15       clientDB.extractIds(emails) >> ch::com
16         >> serverDB::getAttachments >> ch::com
17         >> clientDB::updateAttachments;
18     }
19     else {
20       Choice@Client.ELSE >> ch::select;
21     }
22   }
23 }

```

3.3 A Multitier Programming Implementation with ScalaLoci

We now use ScalaLoci to illustrate the multitier programming approach, implementing the protocol from Figure 3 in Listing 3.

In ScalaLoci, the location of different values is specified through *placement types*. The placement type τ on P represents a value of type τ on a peer P . Developers can freely define the different components, called *peers*, of the distributed system. For instance, in the example, `serverDB` is a `MailServerDB` placed on the `Server` (Line 5) and `clientDB` is a `MailDB` placed on the `Client` (Line 6).

Peers are defined as abstract type members (Lines 2 and 3). Further, peer types express the architectural relation between the different peers by specifying ties between peers, thus supporting generic distributed architectures. Ties statically approximate the run time connections between peers. In the example, we define a *single* tie from client to server (Line 2) and from server to client (Line 3). A single tie expresses the expectation that a single remote instance is always accessible. In the specified architecture, a client connects to a single server and a server program instance handles a single client.

The `updateEmails` method (Line 8) encapsulates the communication logic from Figure 3. It takes the `userId` for identifying the client as input. The implementation diverts control flow to the server using a nested `on[Server].run` expression (Line 10). The `capture` clause transfers both the `timestamp` and the `userId` from the client to the server. Inside the server expression (Line 11), the server queries its local `serverDB` database to extract all e-mails belonging to the `userId` of the client received since its last checkout (indicated by the `timestamp`). The result of the server-side expression is returned to the client using `asLocal` (Line 12).

In ScalaLoci, accessing remote values via the `asLocal` marker creates a local representation of the remote value by transmitting it over the network. Usually, such local representation uses a future, accounting for network delay and potential communication failure. Futures represent values that become available in the future or produce an error. In the example, however, we picked a different transmission scheme where values are transmitted synchronously. ScalaLoci

23:10 Multiparty Languages: The Choreographic and Multitier Cases

■ **Listing 3** ScalaLoci implementation for the context-aware e-mail fetching example.

```
1 @multitier object EmailSystem {
2   @peer type Client <: { type Tie <: Single[Server] }
3   @peer type Server <: { type Tie <: Single[Client] }
4
5   private val serverDB: MailServerDB on Server = ...
6   private val clientDB: MailDB on Client = ...
7
8   def updateEmails(userId: UserId): Unit on Client = on[Client] {
9     val timestamp: Timestamp = clientDB.latestCheckout
10    val emails: List[Email] = on[Server].run.capture(userId, timestamp) {
11      serverDB.since(userId, timestamp)
12    }.asLocal
13
14    clientDB.update(emails)
15    if (ClientLib.isOnFlatRate) {
16      val ids = clientDB.extractIds(emails)
17      clientDB.updateAttachments(
18        on[Server].run.capture(ids) { serverDB.getAttachments(ids) }.asLocal)
19    }
20  }
21 }
```

359 allows developers to choose among different such *transmitters*.

360 The client then uses the received list of `emails` to update its local `clientDB` database (Line 14).
361 Lines 15–19 implement the optional part of the communication logic from Figure 3. If the client
362 is currently using a flat-rate connection—as indicated by the static `ClientLib.isOnFlatRate`
363 method—the client initiates a second server-side computations using `on[Server].run` (Line 18).
364 The client transfers the IDs of the e-mails (retrieved through `extractIds(emails)`)—whose
365 attachments to receive—to the server, which extracts the attachments from its `serverDB`
366 database and returns them to the client, which then updates its local `clientDB` with the
367 received attachments (Line 17).

368 4 Mini Choreographic and Multitier Languages

369 We now introduce Mini Choral and Mini ScalaLoci, minimal languages that omit most features
370 of their reference counterparts that are irrelevant to our study (like generics and inheritance).
371 This allows us to focus on the distinctive traits that characterise the choreographic and
372 multitier approaches, respectively. The minimality of the two languages is instrumental to
373 highlight their distinguishing features here and to focus on the salient points that define their
374 reciprocal translations in Section 5. Next, we present the grammar of the two languages and
375 briefly describe the components that mark them respectively as choreographic and multitier
376 languages.

377 4.1 Mini Choral

378 Listing 4 displays the grammar of Mini Choral. *C* ranges over class declarations, *Channel*
379 ranges over channel declarations, *Field* ranges over class fields, *Method* ranges over method
380 definitions, *Type* ranges over type expressions, and *Exp* ranges over expression terms. The
381 metavariable *id* ranges over both class names, fields, and variables. We use **A**, **B**, **C** to range
382 over role names. Here and in the reminder of the paper, we use *overlines* to denote sequences
383 of terms of the same sort and we denote concatenation of sequences using a comma.

■ **Listing 4** Syntax of Mini Choral

Mini Choral	C	::=	class $id@(\bar{A})\{ \overline{Channel} \overline{Field} \overline{Method} \}$
Type Expression	$Type$::=	$id@(\bar{A})$
Channels	$Channel$::=	$DiChannel@(\bar{A}, \bar{B}) \text{ ch_A_B}$
Field	$Field$::=	$Type \text{ id}$
Method Definition	$Method$::=	$Type \text{ id}(\overline{Type \text{ id}})\{ \text{return} \text{ Exp} \}$
Expression	Exp	::=	$id \mid Exp.id \mid Exp.id(\overline{Exp}) \mid \text{new } id@(\bar{A})(\overline{Exp})$ $\mid \text{lit}@(\bar{A}) \mid \text{if} (Exp) \{ Exp \} \text{ else } \{ Exp \} \mid Exp; Exp$ $\mid \text{ch_A_B.com}(Exp) \mid \text{ch_A_B.select}(Exp)$

384 The class declaration C defines its name id , its owner roles \bar{A} within the $@(\dots)$ clause,
 385 the topology of directed channels available between roles in $\overline{Channel}$, its field declarations
 386 \overline{Field} , and its suite of method definitions \overline{Method} .

387 In Mini Choral, we decided to focus on describing *data flow* and to limit Choral’s
 388 expressivity regarding *data distribution*. That is, we allow only the declared **class** to be
 389 distributed at multiple roles, while variables belong to only one role, with the exception of
 390 *Channels*, which specify the network topology as a set of objects located (and able to transfer
 391 single-role objects) between two roles. Specifically, Mini Choral supports only one-way
 392 channels (drawn from Choral and called *DiChannels*) of the shape $DiChannel@(\bar{A}, \bar{B}) \text{ ch_A_B}$ —
 393 with \bar{A} and \bar{B} roles of the enclosing class. In this work, the loss of expressiveness of the Mini
 394 variant with respect to Choral—which supports the definition of multi-role classes/fields
 395 without the above limitations—lends itself to simplify the algorithms in our translation in
 396 Section 5 without losing generality on the choreographic approach. In the general case, Choral
 397 can express arbitrary channel topologies and user-defined implementations of communications
 398 semantics (e.g., asymmetric channels or bidirectional symmetric channels) [16]—whereas
 399 most choreographic languages assume a complete topology of channels between all roles in a
 400 choreography with a fixed communication semantics [4, 11].

401 Following the considerations above, we restrict type expressions $Type$ to define variables
 402 located at one role $id@(\bar{A})$. This is reflected in the definition of *Fields* but also in method
 403 definitions, where we additionally assume the return type $Type$ and the types of arguments
 404 $\overline{Type \text{ id}}$ to be located at the same role. The body of the method is the single statement
 405 **return** Exp . Regarding expressions, we focus our description on the relevant, non-standard
 406 elements: object creation $\text{new } id@(\bar{A})(\overline{Exp})$ happens for classes at only one role and literals
 407 $\text{lit}@(\bar{A})$ (integers, strings, etc.) are always located at one role. In Exp , we use $Exp; Exp$ to
 408 represent a block which evaluates the expression on the left, discards its value, and returns
 409 the evaluation of the expression on the right.

410 Although already captured by the grammar, we include channel invocations of the shape
 411 $\text{ch_A_B.com}(Exp)$ and $\text{ch_A_B.select}(Exp)$ to highlight their relevance in the language.
 412 *DirectChannels* support both methods **com**, meant to transfer data between two roles, and
 413 **select**, used to solve knowledge-of-choice challenges in conditionals (that is, informing a role
 414 of a local choice made by another role, e.g., by using a conditional) [16]. When using **selects**,
 415 we assume that the compiler provides us with a **Choice** **enum** class at one role, with a **THEN** and
 416 **ELSE** inhabitants (as presented at Line 1 in Listing 2).

417 4.1.1 Example: Mini Choral Expressiveness

418 We conclude the presentation of our minimal choreographic language by illustrating its
 419 expressiveness with respect to its reference Choral language with an implementation of the

23:12 Multiparty Languages: The Choreographic and Multitier Cases

■ **Listing 5** Mini Choral implementation for the context-aware email fetching example.

```
1 class EmailSystem@(Client, Server) {
2   DirectChannel@(Client, Server) ch_Client_Server
3   DirectChannel@(Server, Client) ch_Server_Client
4
5   MailServerDB@Server serverDB
6   MailDB@Client clientDB
7
8   Unit@Client updateEmails(UserId@Client userId) {
9     return contextAwareUpdate(getEmails(userId, clientDB.lastCheckout()))
10  }
11
12  List@Client getEmails(UserId@Client id, Timestamp@Client ts) {
13    return ch_Server_Client.com(
14      serverDB.since(ch_Client_Server.com(id), ch_Client_Server.com(ts))
15    )
16  }
17
18  Unit@Client contextAwareUpdate(List@Client emails) {
19    clientDB.update(emails);
20    if (ClientLib.isOnFlatRate()) {
21      ch_Client_Server.select(Choice@Client.THEN);
22      clientDB.updateAttachments(
23        ch_Server_Client.com(
24          serverDB.getAttachments(
25            ch_Client_Server.com(clientDB.extractIds(emails))))
26    )
27    else {
28      ch_Client_Server.select(Choice@Client.ELSE); Unit
29    }
30  }
```

420 email-fetching protocol presented in Section 3.2, Listing 2.

421 We report the code of the Mini Choral implementation of the protocol in Figure 3 in
422 Listing 5. In the Listing, the main notable difference with Listing 2 is that, by removing
423 assignments, we rely on method bindings to reuse variables in “subsequent” (;) invocations.
424 Although divided into three sub-methods, we find the `updateEmails` method that invokes the
425 `getEmails` method, which fetches the emails from the `Server` by sending to it the `id` of the
426 user and the timestamp (`ts`) of the last checkout and transmitting back the result of the
427 extraction on the `serverDB`. Notice that the return type of the `getEmails` method omits the
428 definition of the “content” of the list due to the lack of generics. As expected, by omitting
429 generics we also drop support for specifying/checking the correct/expected content of the
430 collection—an orthogonal guarantee with respect to the specification/check of the flow of
431 data among roles. The lack of generics does not hamper the expressiveness of the language
432 to capture the correct movement of the data from the `Server` to the `Client` and vice versa.
433 After obtaining the emails, we can apply method `contextAwareUpdate` which updates the email
434 database of the client and proceeds to conditionally retrieve the attachments of the fetched
435 emails. This is done by informing the `Server` of the choice, via the `select` methods.

436 4.2 Mini ScalaLoci

437 Listing 6 displays the grammar of Mini ScalaLoci. *L* ranges over object declarations, *Peer*
438 ranges over peer declarations, *Field* ranges over class fields, *Method* ranges over method
439 definitions, *Type* ranges over type expressions, *PlacedType* ranges over placement type
440 expressions, *Exp* ranges over expressions, and *PlacedExp* ranges over placed expressions. The

■ **Listing 6** Syntax of Mini ScalaLoci

Mini ScalaLoci	L	::=	<code>@multitier object { \overline{Peer} \overline{Field} \overline{Method} }</code>
Peer	$Peer$::=	<code>@peer type A <: { type Tie <: Any with Single[B] }</code>
Placement Type Expression	$PlacedType$::=	<code>Type on A</code>
Type Expression	$Type$::=	<code>id</code>
Field	$Field$::=	<code>val id : PlacedType</code>
Method Definition	$Method$::=	<code>def id ($id : Type$) : PlacedType = PlacedExp</code>
Placed Expression	$PlacedExp$::=	<code>on[A] { Exp }</code>
Expression	Exp	::=	<code>id Exp.id Exp.id(\overline{Exp}) new id(\overline{Exp}) lit if (Exp) { Exp } else { Exp } Exp; Exp on[A].run.capture(id) { Exp }.asLocal</code>

441 metavariable id ranges over both class names, fields, and variables. We use A , B , C to range
442 over peers.

443 The object declaration L defines its name id , and its peers \overline{A} and topology of directed
444 ties between the peers within the `@peer type A <: { type Tie <: Any with Single[A] }` clauses,
445 its field declarations \overline{Field} , and its method definitions \overline{Method} . Fields associate a placement
446 type expression $PlacedType$ to a variable.

447 Mini ScalaLoci is able to express different *topologies* rather than being restricted to a
448 *fixed client-server* model. This choice remarks the departure taken by ScalaLoci from other
449 multitier models and implementations [8, 9, 34, 35, 36], which assume a fixed client-server
450 or n -tier architecture of an application. Contrarily, in ScalaLoci, the developer defines an
451 arbitrary number of peers and directional ties between them. In contrast to ScalaLoci,
452 Mini ScalaLoci only supports a single connected peer instance per peer type (drawn from
453 ScalaLoci’s `Single` ties) of the shape `@peer type A <: { type Tie <: Single[A] }`—with A and B
454 peers of the enclosing multitier module. In this work, the loss of expressiveness of the Mini
455 variant with respect to ScalaLoci lends itself to simplify the algorithms in our translation in
456 Section 5 with losing generality on the multitier approach.

457 In method definitions, the return type $PlacedType$ specifies a location, which places the
458 computation of the whole method on that peer, whereas the arguments only have types
459 but no placement $id : Type$. The body of the method is a placed expression $PlacedExp$ that
460 specifies the placement of the contained expression Exp . Regarding expressions, we focus our
461 description on the main differences with Choral: In ScalaLoci, we locate expressions rather
462 than data and therefore neither instantiation `new id(\overline{Exp})` nor literals `lit` (integers, strings,
463 etc.) carry placement annotations.

464 Nested remote blocks are encoded by `on[A].run.capture(id) { Exp }.asLocal` expressions,
465 which execute the nested expression on the peer A and returns its result via `asLocal` to
466 the surrounding peer, i.e., switching the current perspective to another peer for evaluating
467 the nested expression. Note that in the Mini variant, we keep the `run`, `capture` and `asLocal`
468 constructs to be close to the complete version of the ScalaLoci language (that is syntactically
469 more flexible and supports optional `capture` clauses and `asLocal` on module-level value
470 bindings).

471 4.2.1 Example: Mini ScalaLoci Expressiveness

472 We show the implementation of the email-fetching example presented in Section 3.3, Listing 3
473 using our minimal multitier language to demonstrate its expressiveness with respect to its
474 reference ScalaLoci language.

475 Listing 5 shows the Mini ScalaLoci implementation of the communication scheme in

■ **Listing 7** Mini ScalaLoci implementation for the context-aware email fetching example.

```

1 @multitier object EmailSystem {
2   @peer type Client <: { type Tie <: Any with Single[Server] }
3   @peer type Server <: { type Tie <: Any with Single[Client] }
4
5   val serverDB: MailServerDB on Server
6   val clientDB: MailDB on Client
7
8   def updateEmails(userId: UserId): Unit on Client = on[Client] {
9     contextAwareUpdate(getEmails(userId, clientDB.lastCheckout()))
10  }
11
12  def getEmails(id: UserId, ts: Timestamp): List on Client = on[Client] {
13    on[Server].run.capture(id, ts) {
14      serverDB.since(
15        on[Client].run.capture(id) { id }.asLocal,
16        on[Client].run.capture(ts) { ts }.asLocal)
17    }.asLocal
18  }
19
20  def contextAwareUpdate(emails: List): Unit on Client = on[Client] {
21    clientDB.update(emails);
22    if (ClientLib.isOnFlatRate()) {
23      clientDB.updateAttachments(on[Server].run.capture(emails) {
24        serverDB.getAttachments(
25          on[Client].run.capture(emails) { clientDB.extractIds(emails) }.asLocal
26        }).asLocal)
27    }
28    else {
29      Unit
30    }
31  }
32 }

```

476 Figure 3. As with Mini Choral, the main notable difference with Listing 3 is that by removing
477 assignments, we rely on method arguments for scoped variable declarations instead. The
478 `updateEmails` method invokes the `getEmails` method, which fetches the emails from the `Server`
479 by sending to it the `id` of the user and the timestamp (`ts`) of the last checkout and transmitting
480 back the result of the extraction on the `serverDB`. Similar to Mini Choral, Mini ScalaLoci
481 also lacks generics, an orthogonal language feature. The lack of generics, however, does not
482 limit the expressiveness of the language to capture the correct topology of the system and
483 communication between the `Server` and the `Client`. After obtaining the emails, we apply
484 method `contextAwareUpdate`, which updates the email database of the client and proceeds to
485 conditionally retrieve the attachments of the fetched emails.

486 5 Choreographies to Multitier, Multitier to Choreographies

487 We now define algorithms that translate programs in a Mini language to the other and
488 vice versa. The reason for defining the following algorithms is to present evidence of the
489 existence of a common root at the foundation of the two approaches. We show that the
490 mechanised procedures for their reciprocal translation are relatively simple. In the remainder
491 of this section, for brevity, we use the names Choral and ScalaLoci to indicate their Mini
492 counterparts. We first present a translation algorithm from a Choral choreography to a
493 ScalaLoci multitier application (Section 5.2). Afterwards, we show a translation algorithm
494 from a ScalaLoci multitier application to a Choral choreography (Section 5.2).

495 **Perspective translation** Multitier and choreographic programming take different perspect-
 496 ives on what parts of the language are annotated with locations. In Choral, all literals are
 497 annotated by the role on which they operate, and the location of operators can be inferred
 498 by the location of their argument. ScalaLoc*i* assigns peers to expressions, which are then
 499 written from the specified peer's perspective.

500 While in simple cases there is a direct correspondence between a value on the role *A* in
 501 Choral (`1@A`) and on a peer *A* in ScalaLoc*i* (`on[A] { 1 }`), the difference is more obvious in
 502 compound expressions (`on[A] { 1 + 2 + 3 }` vs. `1@A + 2@A + 3@A`), where in ScalaLoc*i*, only the
 503 whole expression is annotated but the literals are not, whereas in Choral, only the literals
 504 are annotated while the expression is not.

505 The translation algorithms perform such perspective change by grouping composed literals
 506 on the same Choral role into a ScalaLoc*i* placed expression and, in the opposite direction,
 507 assigning the same Choral role to all literals in a ScalaLoc*i* placed expression.

508 Further, we translate between ScalaLoc*i*'s way of defining peers and their topology as type
 509 members and Choral's way of defining roles as class parameters and their communication
 510 channels as class members.

511 **Communication translation** In ScalaLoc*i* two peers communicate using `asLocal`. Given an
 512 expression *e* on peer *A*, the expression `on[B] { e.asLocal }` describes how peer *B* can access the
 513 value of *e*, implemented as a message with the value of *e* sent from *A* to *B*. In Choral, such
 514 communication is represented by invoking the `com` method of a directional communication
 515 channel, which takes a value on role *A* and returns it on role *B*.

516 The translation algorithms transform `asLocal` in ScalaLoc*i* to an invocation of method
 517 `com` of the appropriate channel in Choral and vice versa.

518 5.1 From Choreographic Programming to Multitier Programming

519 **Choral choreography classes to ScalaLoc*i* multitier objects** Algorithm 1 describes the
 520 translation of Choral choreography classes to ScalaLoc*i* multitier objects. We decompose the
 521 class definition to be transformed into its identifier *id*, the roles \overline{Role} , the channel declarations
 522 $\overline{Channel}$, the field declarations \overline{Field} and the method definitions \overline{Method} .

523 Each Choral role definition is translated to a ScalaLoc*i* peer definition. Each channel
 524 `DiChannel@(A,B) ch_A_B` between two roles is translated to a single tie, e.g., a directed one-to-one
 525 tie, between two peers `@peer type A <: { type Tie <: Single[B] }`.

526 The translation of field definition from Choral to ScalaLoc*i* is straightforward. In Choral,
 527 fields are introduced with a base type and the residing role, followed by the name of the field
 528 `"idname@(idrole) idtype"`. In ScalaLoc*i*, fields are introduced as `"val idname: idtype on idrole"`.
 529 Similarly, method definitions are translated.

530 Finally, the algorithm returns a multitier object with the same name and the translated
 531 definitions as a body.

532 **Choral choreography expressions to ScalaLoc*i* multitier expressions** Algorithm 2 de-
 533 scribes the translation of Choral expressions to ScalaLoc*i*: the algorithm matches on the
 534 different cases of Choral *Exp* terms and transforms each into the corresponding ScalaLoc*i*
 535 code.

536 For sequencing $e_0; e_1$, both e_0 and e_1 are recursively transformed. If both subexpressions
 537 agree on their placement, e.g., $A = B$, the complete sequence is placed on the same peer.
 538 More interestingly, if the subexpressions are placed on different peers, we introduce a nested

■ **Algorithm 1** Translation algorithm from Choral classes to ScalaLoci objects.

```

function choral2loci(class)
  "class id@(Role) { Channel Field Method }" ← class
  decls ← { }
  for T ← Role do
    | ties ← { "Single[B]" | "DiChannel@(A, B) ch_A_B" ∈ Channel ∧ T = A }
    | decls ← decls ∪ { "@peer type T <: { type Tie <: Any with ties }" }
  end
  for "id@(A) idn" ← Field do
    | decls ← decls ∪ { "val id1 : id0 on A" }
  end
  for "id@(A) id (idtn@(A) iden) { e }" ← Method do
    | e' ← choral2loci(e)
    | decls ← decls ∪ { "def id(iden : idtn) : idt on A = {e'}" }
  end
  return "@multitier object id { decls }"
end

```

539 remote block for e'_0 , which executes e'_0 on A and places the overall result of e'_1 on B . For the
 540 remote block we generate a capture clause for all method-local variables that are free in e_0 .

541 The translations for identifiers, literals and instantiation is straightforward, placing the
 542 ScalaLoci expression on the peer according to the role specified in the Choral code. Further,
 543 the case for method invocation is similar since we assume that the receiver of a method
 544 invocation and its arguments are on the same role. This assumption is expressed by the *assert*
 545 statement in the algorithm and holds for every well-typed Mini Choral program. Selection
 546 does not exist in ScalaLoci. Hence, it is removed.

547 The case for branching makes a distinction similar to sequencing of whether the condition
 548 agrees to the branches regarding their placement, e.g., $A = B$. If they agree, the complete
 549 branching is placed on the same peer. Otherwise, we introduce a nested remote block for e'_0 ,
 550 which executes e'_0 on A and returns the result to B where the branches are placed. B then
 551 acts as a coordinator to decide which of the branches to execute.

552 Finally, we translate Choral's channel communication. For a channel from role B to A ,
 553 we generate a ScalaLoci expression, which runs a nested remote block for e' , which executes
 554 e' on B and returns the result to A .

555 5.2 From Multitier Programming to Choreographic Programming

556 **ScalaLoci multitier objects to Choral choreography classes** Algorithm 3 describes the
 557 translation of ScalaLoci multitier objects to Choral choreography classes. We decompose the
 558 multitier object to be transformed into its identifier *id*, the peer and tie declarations \overline{Peer} ,
 559 the field declarations \overline{Field} and the method definitions \overline{Method} .

560 Each ScalaLoci peer definition is translated to Choral role argument and each single tie
 561 between two peers is translated to a **DiChannel** between two peers **@(A,B)**.

562 The translation of field and method definitions from ScalaLoci to Choral is straightforward.

563 Finally, the algorithm returns a Choral class with the same name and the translated
 564 definitions as a body.

■ **Algorithm 2** Translation algorithm from Choral expressions to ScalaLoci expressions.

```

function choral2loci(expr)
  return match expr with
  | case "e0; e1" with
    | "on[A]{ e'0 }" ← choral2loci(e0)
    | "on[B]{ e'1 }" ← choral2loci(e1)
    | captures ← freeVars(e0) ∩ currentMethodArguments
    | if A ≠ B then
      | "on[B]{ on[A].run.capture(captures) { e'0 }.asLocal; e'1 }"
    | else
      | "on[B]{ e'0; e'1 }"
    | end
  | case "id" with
    | A ← roleOf(id)
    | "on[A]{ id }"
  | case "lit@A" with
    | "on[A]{ lit }"
  | case "new id@A (ē)" with
    | "on[A]{ e' }" ← choral2loci(e)
    | "on[A]{ new id(e') }"
  | case "e0.id(ē)" with
    | "on[A]{ e'0 }" ← choral2loci(e0)
    | "on[B]{ e' }" ← choral2loci(e)
    | assert A = B // receiver and arguments have the same role
    | "on[A]{ e'0.id(e') }"
  | case "ch.select(e)" with
    | "Unit"
  | case "if (e0) { e1 } else { e2 }" with
    | "on[A]{ e'0 }" ← choral2loci(e0)
    | "on[B]{ e'1 }" ← choral2loci(e1)
    | "on[C]{ e'2 }" ← choral2loci(e2)
    | captures ← freeVars(e0) ∩ currentMethodArguments
    | assert B = C // branches have the same role
    | if A ≠ B then
      | "on[B]{ if (on[A].run.capture(captures) { e'0 }.asLocal) { e'1 } else { e'2 } }"
    | else
      | "on[B]{ if (e'0) { e'1 } else { e'2 } }"
    | end
  | case "ch_B_A.com(e)" with
    | "on[B]{ e' }" ← choral2loci(e)
    | captures ← freeVars(e) ∩ currentMethodArguments
    | "on[A]{ on[B].run.capture(captures) { e' }.asLocal }"
  end
end
end

```

565 **ScalaLoci multitier expressions to Choral choreography expressions** Algorithm 4 de-
 566 scribes the translation of ScalaLoci expressions to Choral expressions. The algorithm matches
 567 on the different cases of ScalaLoci *Expr* terms and transforms each of them into the corres-
 568 ponding ScalaLoci code.

569 The translations for sequencing, identifiers, literals, instantiation and method invocation
 570 is straightforward, recursively transforming each subexpression.

571 In the case for branching, the translation needs to synthesise select expressions to
 572 implement knowledge of choice (recall Section 3.2). Hence, we collect all peers used in the

■ **Algorithm 3** Translation algorithm from ScalaLoci objects to Choral classes.

```

function loci2choral(module)
  "@multitier object id {  $\overline{Peer}$   $\overline{Field}$   $\overline{Method}$  }"  $\leftarrow$  module
  decls  $\leftarrow$  { }
  roles  $\leftarrow$  { }
  for "@peer type A <: { type Tie <: Any with ties }"  $\leftarrow$  Peer do
    | roles  $\leftarrow$  roles  $\cup$  { A }
    | for "Single[B]"  $\leftarrow$  ties do
      | | decls  $\leftarrow$  decls  $\cup$  { "DiChannel@(A,B) ch_A_B" }
      | end
    | end
  | for "val id1: id0 on A"  $\leftarrow$  Field do
    | | decls  $\leftarrow$  decls  $\cup$  { "id0@(A) id1" }
    | end
  | for "def id(iden: idtn): idt on A = { e }"  $\leftarrow$  Method do
    | | e'  $\leftarrow$  loci2choral(e)
    | | decls  $\leftarrow$  decls  $\cup$  { "idt@(A) id(idtn@(A) iden) { e' }" }
    | end
  | return "class id@(roles) {  $\overline{decls}$  }"
end

```

573 branches and create select statements for all channels between those peers for both branches.

574 Finally, we translate ScalaLoci's nested remote blocks. For a remote expression placed
 575 on *A* that executes *e* on *B*, we generate a Choral channel communication that transfers the
 576 value of *e* from *B* to *A*.

577 6 A Unified Perspective

578 Although choreographic and multitier programming evolved in dissimilar ways, their cores—
 579 represented by our two Mini languages—are close enough to let us define in Section 5
 580 straightforward translation algorithms in both directions and show the core features of both
 581 approaches isomorphic.

582 Besides the more abstract purpose to present evidence of the closeness of the two
 583 approaches, our translation algorithms are also directly useful in practice. Translating
 584 Choral to ScalaLoci code enables the reuse of ScalaLoci's middleware for Choral. In general,
 585 translating to multitier programs is interesting because we can leverage the possibility of
 586 compiling to different technologies.

587 Translating ScalaLoci to Choral code enables synthesising the choreography of the
 588 multitier program. Making the protocol manifest supports both manually checking what
 589 communications take place as well as automatic analyses (e.g., security).

590 We believe that both the multitier and choreographic research areas can greatly benefit
 591 from cross-fertilisation and transfer of concepts already developed in one but lacking in the
 592 other. As a glimpse of this fact, we dedicate Section 6.1 to describe some advanced features
 593 present in only one of the two languages (Choral, ScalaLoci) and outline how they could be
 594 integrated into the other in the future. We conclude this section by widening our scope on
 595 the category of multiparty language in Section 6.2. We give an (incomplete) overview on
 596 other languages that are neither multitier nor choreographic but share common traits that
 597 can classify them as multiparty ones. We consider those languages valuable additions to the

■ **Algorithm 4** Translation algorithm from ScalaLoci expressions to Choral expressions.

```

function loci2choral(expr)
  return match expr with
  | case "on[A]{ e0; e1 }" with
    | e'0 ← loci2choral("on[A]{ e0 }")
    | e'1 ← loci2choral("on[A]{ e1 }")
    | "e'0; e'1"
  | case "on[A]{ id }" with "id"
  | case "on[A]{ lit }" with
    | "lit@A"
  | case "on[A]{ new id( $\bar{e}$ ) }" with
    | e' ← loci2choral("on[A]{ e }")
    | "new id@A(e')"
  | case "on[A]{ e0.id( $\bar{e}$ ) }" with
    | e'0 ← loci2choral("on[A]{ e0 }")
    | e' ← loci2choral("on[A]{ e }")
    | "e'0.id(e')"
  | case "on[A]{ if (e0) { e1 } else { e2 } }" with
    | e'0 ← loci2choral("on[A]{ e0 }")
    | e'1 ← loci2choral("on[A]{ e1 }")
    | e'2 ← loci2choral("on[A]{ e2 }")
    | peers ← peersIn(e'1) ∪ peersIn(e'2)
    | channels ← { "ch_A_B" | B ∈ peers ∧ A has tie to B }
    | thenSelects ← { "c.select(Choice@A.THEN)" | c ∈ channels }
    | elseSelects ← { "c.select(Choice@A.ELSE)" | c ∈ channels }
    | "if ( e'0 ) { thenSelects; e'1 } else { elseSelects; e'2 }"
  | case "on[A]{ on[B].run.capture(captures) { e }.asLocal }" with
    | e' ← loci2choral("on[B]{ e }")
    | "ch_B_A.com(e')"
  end
end

```

Feature	Choral	ScalaLoci
Distributed Data Structures	✓	×
Dynamic Topologies	×	✓
Higher-Order Composition	✓	×
Races	—	✓
Fault tolerance	✓	✓
Asynchrony	✓	✓

■ **Table 1** Overview of the feature comparison of choreographic and multitier programming.

598 multiparty category and subject of future research akin to this work.

599 6.1 Feature Comparison

600 We now discuss a few features that are important for concurrent and distributed programming.
 601 Our discussion is summarised in Table 1, which shows which features are present in Choral and
 602 ScalaLoci, respectively (the — means partial presence, explained in the relevant paragraph
 603 where we discuss the feature). The first four features have evolved separately and give

23:20 Multiparty Languages: The Choreographic and Multitier Cases

604 potential for cross-fertilisation, whereas the last two are important features that have been
605 dealt in both worlds (yet separately).

606 Distributed Data Structures

607 The `@(R1, ..., Rn)` type notation supported in Choral specifies the distribution of classes
608 and objects over roles. This is true also without taking into account communication. As an
609 example, let us consider the `BiPair` class below, which implements an incarnation of a `Pair`
610 class where the two values (referred as `left` and `right`) of the pair belong to different roles.

```
611 class BiPair@(A,B)<L@X, R@Y> {  
612   private L@A left;  
613   private R@B right;  
614   public BiPair(L@A left, R@B right) { this.left = left; this.right = right; }  
615   public L@A left() { return this.left; }  
616   public R@A right() { return this.right; }  
617 }  
618
```

620 As its Java counterpart, also `BiPair` is parametric with respect to its contents: we use
621 parameters `L` and `R` to capture the type of the left and right components of the pair. Then,
622 by specifying that `L` is owned by one role `X` and `R` is owned by another role `Y`, we indicate that
623 the two values in the pair must be at different roles (and they can capture different data
624 types, e.g., `String` and `Integer`). Indeed, adopting the same interpretation of Java generics,
625 Choral interprets role parameter binders so that the first appearance of a parameter is a
626 binder, while subsequent appearances of the same parameter are bound—hence, given that
627 the declaration of type parameters `<...>` limits the scope of the of role parameters `X` and `Y`,
628 we are indicating that they cannot coincide. For completeness, we include in the definition
629 of the `BiPair` class its fields (`left` and `right`, respectively located at `A` and `B`), a constructor,
630 and the traditional accessors.

631 Besides showing the basic feature of inherent distribution supported by the Choral
632 type system, the example of `BiPair` is useful to illustrate that, also without considering
633 communications, Choral offers support in defining programs where the data at some role
634 needs to correlate with data at another, e.g., as in the case of distributed authentication
635 tokens.

636 Similar to Choral, in `ScalaLoc`i, we use parameters `L` and `R` to capture the type of the left
637 and right components of the pair. Corresponding to Choral's roles definition, we define an `A`
638 and a `B` peer type. We then specify that `L` is placed on a peer `A` and `R` is placed on a peer `B`:

```
639 @multitier trait BiPair[L, R] {  
640   @peer type A  
641   @peer type B  
642  
643   val left: L on A  
644   val right: R on B  
645 }  
646  
647
```

648 While we can define distributed data structures similar to Choral, they have to be
649 composed at compile-time limiting their usability because of `ScalaLoc`i's lack of higher-order
650 composition.

651 Dynamic Topologies and Homogenous Behaviours


652 A feature of `ScalaLoc`i that is not covered in its Mini variant is the possibility for peer types
653 to abstract over multiple peer instances of the same type, e.g., a master-worker architecture
654 where a single master can connect to an arbitrary number of homogeneous (i.e., with the

■ Listing 8 Distributed Architectures.

```

1 @multitier object P2P {
2   @peer type Peer <: { type Tie <: Multiple[Peer] }
3 }
4 @multitier object P2PRegistry {
5   @peer type Registry <: { type Tie <: Multiple[Peer] }
6   @peer type Peer <: { type Tie <: Optional[Registry] with Multiple[Peer] }
7 }
8 @multitier object MultiClientServer {
9   @peer type Server <: { type Tie <: Multiple[Client] }
10  @peer type Client <: { type Tie <: Single[Server] with Single[Node] }
11 }
12 @multitier object ClientServer {
13   @peer type Server <: { type Tie <: Single[Client] }
14   @peer type Client <: { type Tie <: Single[Server] with Single[Node] }
15 }
16 @multitier object Ring {
17   @peer type Node <: { type Tie <: Single[Prev] with Single[Next] }
18   @peer type Prev <: Node
19   @peer type Next <: Node
20   @peer type RingNode <: Prev with Next
21 }

```



655 same behaviour) worker nodes. Such a feature also enables dynamic topologies where peers
 656 can join and leave the system at run time. A variable number of peer instances is expressed
 657 in ScalaLoci’s peer specification by not using a `Single` tie but an `Multiple` or an `Optional` tie,
 658 i.e., an arbitrary number or at most one remote peer of a given type can connect, respectively.

659 Listing 8 shows the definitions for different topologies with their iconification on the
 660 right. The `P2P` module defines a `Peer` that can connect to arbitrary many other peers. The
 661 `P2PRegistry` module adds a central registry, to which peers can connect. The `MultiClient-`
 662 `Server` module defines a client that is always connected to a single server, while the server
 663 can handle multiple clients simultaneously. The `ClientServer` module specifies a server that
 664 always handles a single client instance. For the `Ring` module, we define a `Prev` and a `Next`
 665 peer. A `RingNode` itself is both a predecessor and a successor. All `Node` peers have a single
 666 tie to their predecessor and a single tie to their successor.

667 ScalaLoci allows to abstract over different peer instances of the same type and uniformly
 668 receive values from multiple connected remote peers, `asLocalFromAll` returns a sequence that
 669 contains the remote values from the different peers. Yet, a specific peer instance `client`
 670 can be selected via `on(client).run { ... }.asLocal` (using the `client` value referencing a peer
 671 instance) instead of `on[Client].run { ... }.asLocal` (using the `Client` peer type). The handlers
 672 `remote[Client].join foreach { ... }` and `remote[Client].leave foreach { ... }` can be used to
 673 react to dynamic changes in the topology of the running multitier system.

674 Denièlou and Yoshida [13] developed a theory for choreographies with homogeneous roles
 675 and dynamic topologies by allowing choreographies to be parametrised (also) in collections of
 676 roles. Plans for supporting for this feature in Choral are discussed in [16, §7]. In this extension,
 677 prefixing a role parameter declaration with `*`, as in `*Clients`, specifies that this is a collection
 678 of roles. Types are extended with products indexed over collections of role using a syntax
 679 similar to Java `for`-each blocks. For instance, the type `forall(Client: Clients) String@Client`
 680 represents a “tuple” with a `String` for each role in the collection `Clients`. We can write a
 681 scatter-gather channel over a star topology (cf. `MultiClientServer`) as follows:

```

682
6831 abstract class StarChannel@(Server,*Clients) {
6842   forall(Client : Clients) {SymChannel@(Server,Client)} star;
6853   forall(Client : Clients) {String@Client} scatter(String@Server m);
6864   String@Server gather( forall(Client : Clients) {String@Client} ms);
6875 }

```

23:22 Multiparty Languages: The Choreographic and Multitier Cases

688

689 Method `gather` of `StarChannel` is then translated to ScalaLoci’s primitive `asLocalFromAll` and
690 vice versa. A further extension discussed in [16, §7] is the introduction of existential quan-
691 tification over roles in role collections. For instance, `with(Client: Clients) String@(Client)`
692 represents a string at some role in the collection `Clients`. We can extend the example above
693 to support any-cast communication as follows:

```
694 abstract class StarChannel@(Server,*Clients) {  
695   /* ... */  
696   with(Client : Clients) {String@(Client)} any( String@Server m );  
697   String@Server any( with(Client : Clients) {String@(Client)} m );  
698 }  
699
```

701 Method `any` of `StarChannel` is then translated to ScalaLoci’s `on(c).run { ... }` and vice
702 versa.

703 Higher-Order and First-Class Multiparty Programs

704 We classify “higher-order” a multiparty language where multiparty components (objects,
705 functions) are values that can be passed as arguments.

706 Choral is higher-order because methods can accept choreographic objects with multiple
707 roles as parameters. In Choral, channels are one of the most basic examples of the usage of
708 the higher-order feature. For example, we can pass a `DiChannel` as an argument.

```
709 class MyClass@(A,B){  
710   void passValue( DiChannel@(A,B) ch ){  
711     ch.com< Integer >( 5@B );  
712   }  
713 }  
714 }  
715
```

716 In the example, the method `passValue` takes as input the choreographic object `DiChannel`
717 and, by invoking its `com` method, we execute the protocol needed to send the data (`5@B`)
718 between the two roles.

719 ScalaLoci does not support higher-order composition (no multitier objects as values or
720 dynamic multitier object storage) but at least supports statically-composed modules [39].
721 The following snippet shows the declaration of a `ClientServer` multitier module that is
722 parameterised over a `Client` and a `Server` peer. The module uses the monitoring functionality
723 provided by the `Monitoring` multitier module, which is parameterised over a `Monitor` and
724 a `Monitored` peer. The `Monitoring` module is instantiated by `mon` inside `ClientServer`. The
725 `ClientServer` module identifies the `Client` peer with the `Monitored` peer and the `Server` peer
726 with the `Monitor` peer and defines their ties accordingly:

```
727 @multitier trait Monitoring {  
728   @peer type Monitor { type Tie <: Single[Monitored] }  
729   @peer type Monitored { type Tie <: Single[Monitor] }  
730 }  
731  
732 @multitier object ClientServer {  
733   @multitier object mon extends Monitoring  
734  
735   @peer type Client <: mon.Monitored { type Tie <: Single[mon.Monitor] with Single[Server] }  
736   @peer type Server <: mon.Monitor { type Tie <: Single[mon.Monitored] with Single[Client] }  
737 }  
738
```


740 **Races**

741 Both Choral and ScalaLoci support programs with races to some degree. We distinguish two
742 prototypical scenarios: races among producers and races among consumers.

743 To program a race among multiple producers in ScalaLoci, we can simply retrieve the
744 values from all remote producers via `asLocalFromAll` and pick the first one that becomes
745 available via `Future.firstCompletedOf` as shown in the example below:

```
746 Future.firstCompletedOf(  
747   on[Producer].run { generateValue() }.asLocalFromAll map {  
748     case (producerPeerInstance, value) => value map { (producerPeerInstance, _) }  
749   })  
750 }
```

752 It is not possible to program a race among multiple consumers in ScalaLoci.

753 In Choral, it is possible to implement protocols with races among producers and among
754 consumers provided their number is statically fixed. For instance, below is the type for a
755 choreography where two producers race to send a message to a consumer.

```
756 interface ProducerRace(Producer1,Producer2,Consumer) {  
757   Message@Consumer run(Message@Producer1 m1, Message@Producer2 m2);  
758 }  
759 }
```

761 The constraint that the number of roles must be statically fixed is related to the inability of
762 Choral to capture dynamic topologies and, as discussed above, is solved by adding collections
763 of roles to the language. In the case of consumer races, another limitation is that the Choral
764 type system is not powerful enough express (and enforce) their presence. Consider a situation
765 where two consumers race to receive a message from a single producer. In Choral, this
766 protocol can implement the following interface:

```
767 interface ConsumerRace(Producer,Consumer1,Consumer2) {  
768   BiPair@(Consumer1,Consumer2)<Optional<Message>,Optional<Message>> run(Message@Producer m);  
769 }  
770 }
```

772 However, the return type of `run` does not guarantee that exactly one consumer receives
773 the message: implementations that deliver the message to both or neither respect the type.
774 As discussed in [16, §7], we can write a precise type if we extend Choral with existential
775 quantification over roles as shown in the example below.

```
776 interface ConsumerRace@(Producer, Consumer1, Consumer2) {  
777   with (C : [Consumer1, Consumer2]){ Message@C } run(Message@Producer m);  
778 }  
779 }
```

781 **Fault Tolerance**

782 In ScalaLoci, remote values whose computation or transmission to the local peer instance
783 fail result in a future that is completed with a failure value. Thus, user code can detect a
784 failed remote access and decide how to react appropriately. Failed futures can be handled
785 using the usual operators on futures, e.g., `recover`:

```
786 on[Client].run { generateValue() }.asLocal recover { case _ => generateOtherValue() }  
787 }
```

789 Like other aspects of communication in Choral, the language does not commit to specific
790 interaction patterns: programmers can implement their own strategies e.g. using exceptions
791 or returning errors.

792 **Asynchrony**

793 For the sake of exposition, we presented multiparty programs using communication APIs
 794 as if they were blocking and designed the Mini variants of both Choral and ScalaLoci
 795 as synchronous. ScalaLoci promotes an asynchronous approach: the preferred variant of
 796 accessing remote values via `asLocal` in ScalaLoci creates a future to account for network delay
 797 and potential communication failure. On the other hand, Choral is agnostic with regards to
 798 communication models: programmers can import libraries of channels or implement their
 799 own. For instance, a communication model similar to ScalaLoci’s `asLocal` is offered by the
 800 following interface:

```
801 interface AsyncDiChannel@(<Sender, Receiver><T@X> {
802   <S@Y extends T@Y> Future@Receiver<S> com(Promise@Sender<S>);
803 }
804
805 }
```

806 **6.2 Other Multiparty Languages**

807 For the future we envision further cross-fertilisation between multiparty languages, and that
 808 the class of multiparty languages might get larger. We mention a few approaches that might
 809 contribute to this.

810 Software architectures [14, 33] are about organising software systems into well-studied
 811 patterns that comprise components and their connections organised in a certain configura-
 812 tion. Architecture description languages (ADL) [26] specify software architectures and the
 813 constraints among the architecture components. Different from choreographic and multitier
 814 programming, ADLs usually specification languages separate from the implementation. An
 815 exception is ArchJava [1] which support specifying a software architecture and enforcing
 816 its constraints together with the implementation. Regarding cross-fertilisation, ADLs come
 817 equipped with powerful analysis, code synthesis, and runtime-support tools as well as model
 818 checkers, which can be also used in multitier and choreographic scenarios to enforce different
 819 aspects of correctness.

820 Partitioned global address space languages (PGAS) [12] are often used in the domain
 821 of high-performance computing. The main abstraction is a global memory address space
 822 where logical partitions are assigned to processes to maximize data locality. X10 [7] features
 823 explicit fork/join operations and provides a sophisticated dependent type system [6] to model
 824 the *place* (the heap partition) a reference points to. PGAS languages, similar to multitier
 825 and choreographic languages reduce the boundaries between hosts in a distributed system.

826 **7 Conclusion**

827 Choreographic and multitier languages have developed independently, leading to a number
 828 of research achievement carried out within two vibrant but separate research communities [2,
 829 28, 40]. In this paper, we discussed the fundamental nature of the programming paradigms
 830 based on these languages, isolating the core difference between them. We then showed that,
 831 under the cover of syntactic variance, the two approaches are similar enough to be related
 832 and to reason about potential cross-fertilisation. Our observations offer a platform for future
 833 joint work between the respective communities.

834 **References**

- 835 1 Jonathan Aldrich, Craig Chambers, and David Notkin. ArchJava: Connecting software
 836 architecture to implementation. In *Proceedings of the 24th International Conference on*

- 837 *Software Engineering*, ICSE '02, pages 187–197, New York, NY, USA, 2002. ACM. doi:
838 10.1145/581339.581365.
- 839 2 Davide Ancona, Viviana Bono, Mario Bravetti, Joana Campos, Giuseppe Castagna, Pierre-
840 Malo Deniérou, Simon J. Gay, Nils Gesbert, Elena Giachino, Raymond Hu, Einar Broch
841 Johnsen, Francisco Martins, Viviana Mascardi, Fabrizio Montesi, Romyana Neykova, Nich-
842 olas Ng, Luca Padovani, Vasco T. Vasconcelos, and Nobuko Yoshida. Behavioral types in
843 programming languages. *Foundations and Trends in Programming Languages*, 3(2-3):95–230,
844 2016.
- 845 3 Marco Carbone, Kohei Honda, and Nobuko Yoshida. Structured communication-centered
846 programming for web services. *ACM Trans. Program. Lang. Syst.*, 34(2):8:1–8:78, 2012.
847 doi:10.1145/2220365.2220367.
- 848 4 Marco Carbone and Fabrizio Montesi. Deadlock-freedom-by-design: multiparty asynchronous
849 global programming. In Roberto Giacobazzi and Radhia Cousot, editors, *The 40th Annual ACM*
850 *SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '13, Rome,*
851 *Italy - January 23 - 25, 2013*, pages 263–274. ACM, 2013. doi:10.1145/2429069.2429101.
- 852 5 Luca Cardelli and Andrew D. Gordon. Mobile ambients. In Maurice Nivat, editor, *Foundations*
853 *of Software Science and Computation Structure, First International Conference, FoSSaCS'98,*
854 *Held as Part of the European Joint Conferences on the Theory and Practice of Software,*
855 *ETAPS'98, Lisbon, Portugal, March 28 - April 4, 1998, Proceedings*, volume 1378 of *Lecture*
856 *Notes in Computer Science*, pages 140–155. Springer, 1998. doi:10.1007/BFb0053547.
- 857 6 Satish Chandra, Vijay Saraswat, Vivek Sarkar, and Rastislav Bodik. Type inference for locality
858 analysis of distributed data structures. In *Proceedings of the 13th ACM SIGPLAN Symposium*
859 *on Principles and Practice of Parallel Programming, PPOPP '08*, pages 11–22, New York, NY,
860 USA, 2008. ACM. doi:10.1145/1345206.1345211.
- 861 7 Philippe Charles, Christian Grothoff, Vijay Saraswat, Christopher Donawa, Allan Kielstra,
862 Kemal Ebcioglu, Christoph von Praun, and Vivek Sarkar. X10: An object-oriented approach to
863 non-uniform cluster computing. In *Proceedings of the 20th Annual ACM SIGPLAN Conference*
864 *on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA '05*, pages
865 519–538, New York, NY, USA, 2005. ACM. doi:10.1145/1094811.1094852.
- 866 8 Ezra Cooper, Sam Lindley, Philip Wadler, and Jeremy Yallop. Links: Web programming
867 without tiers. In Frank S. de Boer, Marcello M. Bonsangue, Susanne Graf, and Willem P.
868 de Roeper, editors, *Formal Methods for Components and Objects, 5th International Symposium,*
869 *FMCO 2006, Amsterdam, The Netherlands, November 7-10, 2006, Revised Lectures*, volume
870 4709 of *Lecture Notes in Computer Science*, pages 266–296. Springer, 2006. doi:10.1007/
871 978-3-540-74792-5_12.
- 872 9 Ezra E. K. Cooper and Philip Wadler. The RPC calculus. In *Proceedings of the 11th ACM*
873 *SIGPLAN Conference on Principles and Practice of Declarative Programming, PPDP '09*,
874 pages 231–242, New York, NY, USA, 2009. ACM. doi:10.1145/1599410.1599439.
- 875 10 Luís Cruz-Filipe and Fabrizio Montesi. Choreographies in practice. In Elvira Albert and Ivan
876 Lanese, editors, *Formal Techniques for Distributed Objects, Components, and Systems - 36th*
877 *IFIP WG 6.1 International Conference, FORTE 2016, Held as Part of the 11th International*
878 *Federated Conference on Distributed Computing Techniques, DisCoTec 2016, Heraklion, Crete,*
879 *Greece, June 6-9, 2016, Proceedings*, volume 9688 of *Lecture Notes in Computer Science*, pages
880 114–123. Springer, 2016. doi:10.1007/978-3-319-39570-8_8.
- 881 11 Mila Dalla Preda, Maurizio Gabbriellini, Saverio Giallorenzo, Ivan Lanese, and Jacopo Mauro.
882 Dynamic choreographies: Theory and implementation. *Logical Methods in Computer Science*,
883 13(2), 2017. doi:10.23638/LMCS-13(2:1)2017.
- 884 12 Mattias De Wael, Stefan Marr, Bruno De Fraine, Tom Van Cutsem, and Wolfgang De Meuter.
885 Partitioned global address space languages. *ACM Computing Surveys*, 47(4), May 2015.
886 doi:10.1145/2716320.
- 887 13 Pierre-Malo Deniérou and Nobuko Yoshida. Dynamic multirole session types. In Thomas Ball
888 and Mooly Sagiv, editors, *Proceedings of the 38th ACM SIGPLAN-SIGACT Symposium on*

- 889 *Principles of Programming Languages, POPL 2011, Austin, TX, USA, January 26-28, 2011,*
890 pages 435–446. ACM, 2011. doi:10.1145/1926385.1926435.
- 891 **14** David Garlan and Mary Shaw. An introduction to software architecture. Technical report,
892 Pittsburgh, PA, USA, 1994. Accessed 2020-05-05. URL: [http://www.cs.cmu.edu/afs/cs/
893 project/vit/ftp/pdf/intro_softarch.pdf](http://www.cs.cmu.edu/afs/cs/project/vit/ftp/pdf/intro_softarch.pdf).
- 894 **15** Saverio Giallorenzo, Ivan Lanese, and Daniel Russo. Chip: A choreographic integration
895 process. In *On the Move to Meaningful Internet Systems. OTM 2018 Conferences - Confederated
896 International Conferences: CoopIS, C&TC, and ODBASE 2018, Valletta, Malta, October 22-26,
897 2018, Proceedings, Part II*, pages 22–40. Springer, 2018. doi:10.1007/978-3-030-02671-4_2.
- 898 **16** Saverio Giallorenzo, Fabrizio Montesi, and Marco Peressotti. Choreographies as objects. *CoRR*,
899 abs/2005.09520, 2020. URL: <https://arxiv.org/abs/2005.09520>, arXiv:2005.09520.
- 900 **17** Carl Hewitt, Peter Bishop, and Richard Steiger. A universal modular Actor formalism for
901 artificial intelligence. In *Proceedings of the 3rd International Joint Conference on Artificial
902 Intelligence, IJCAI '73*, pages 235–245, San Francisco, CA, USA, 1973. Morgan Kaufmann
903 Publishers Inc. Accessed 2020-05-05. URL: [http://ijcai.org/Proceedings/73/Papers/
904 027B.pdf](http://ijcai.org/Proceedings/73/Papers/027B.pdf).
- 905 **18** Kohei Honda, Aybek Mukhamedov, Gary Brown, Tzu-Chun Chen, and Nobuko Yoshida.
906 Scribbling interactions with a formal foundation. In Raja Natarajan and Adegboyega K.
907 Ojo, editors, *Distributed Computing and Internet Technology - 7th International Conference,
908 ICDCIT 2011, Bhubaneswar, India, February 9-12, 2011. Proceedings*, volume 6536 of *Lecture
909 Notes in Computer Science*, pages 55–75. Springer, 2011. doi:10.1007/978-3-642-19056-8\
910 _4.
- 911 **19** Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty asynchronous session types.
912 *J. ACM*, 63(1):9, 2016. Also: *POPL*, pages 273–284, 2008. URL: [http://doi.acm.org/10.
913 1145/2827695](http://doi.acm.org/10.1145/2827695), doi:10.1145/2827695.
- 914 **20** Intl. Telecommunication Union. Recommendation Z.120: Message Sequence Chart, 1996.
- 915 **21** Tanakorn Leesatapornwongsa, Jeffrey F. Lukman, Shan Lu, and Haryadi S. Gunawi. TaxDC:
916 A taxonomy of non-deterministic concurrency bugs in datacenter distributed systems. In *Proc.
917 of ASPLOS*, pages 517–530, 2016.
- 918 **22** Alberto Lluch-Lafuente, Flemming Nielson, and Hanne Riis Nielson. Discretionary information
919 flow control for interaction-oriented specifications. In *Logic, Rewriting, and Concurrency*,
920 volume 9200 of *Lecture Notes in Computer Science*, pages 427–450. Springer, 2015.
- 921 **23** Hugo A. López and Kai Heussen. Choreographing cyber-physical distributed control systems
922 for the energy sector. In *SAC*, pages 437–443. ACM, 2017.
- 923 **24** Hugo A. López, Flemming Nielson, and Hanne Riis Nielson. Enforcing availability in failure-
924 aware communicating systems. In *FORTE*, volume 9688 of *Lecture Notes in Computer Science*,
925 pages 195–211. Springer, 2016.
- 926 **25** Shan Lu, Soyeon Park, Eunsoo Seo, and Yuanyuan Zhou. Learning from mistakes: a com-
927 prehensive study on real world concurrency bug characteristics. In *Proc. of ASPLOS*, pages
928 329–339, 2008.
- 929 **26** Nenad Medvidovic and Richard N. Taylor. A classification and comparison framework for
930 software architecture description languages. *IEEE Transactions on Software Engineering*,
931 26(1):70–93, January 2000. doi:10.1109/32.825767.
- 932 **27** Robin Milner. *A Calculus of Communicating Systems*, volume 92 of *Lecture Notes in Computer
933 Science*. Springer, 1980. doi:10.1007/3-540-10235-3.
- 934 **28** Fabrizio Montesi. *Choreographic Programming*. Ph.D. Thesis, IT University of Copenhagen,
935 2013. http://www.fabriziomontesi.com/files/choreographic_programming.pdf.
- 936 **29** Tom Murphy VII, Karl Cray, Robert Harper, and Frank Pfenning. A symmetric modal
937 lambda calculus for distributed computing. In *19th IEEE Symposium on Logic in Computer
938 Science (LICS 2004), 14-17 July 2004, Turku, Finland, Proceedings*, pages 286–295. IEEE
939 Computer Society, 2004. doi:10.1109/LICS.2004.1319623.

- 940 30 Roger M. Needham and Michael D. Schroeder. Using encryption for authentication in large
941 networks of computers. *Commun. ACM*, 21(12):993–999, 1978.
- 942 31 Object Management Group. Business Process Model and Notation.
943 <http://www.omg.org/spec/BPMN/2.0/>, 2011.
- 944 32 Peter W. O’Hearn. Experience developing and deploying concurrency analysis at facebook. In
945 Andreas Podelski, editor, *Static Analysis - 25th International Symposium, SAS 2018, Freiburg,*
946 *Germany, August 29-31, 2018, Proceedings*, volume 11002 of *Lecture Notes in Computer*
947 *Science*, pages 56–70. Springer, 2018. doi:10.1007/978-3-319-99725-4_5.
- 948 33 Dewayne E. Perry and Alexander L. Wolf. Foundations for the study of software architecture.
949 *ACM SIGSOFT Software Engineering Notes*, 17(4):40–52, October 1992. doi:10.1145/141874.
950 141884.
- 951 34 Gabriel Radanne, Jérôme Vouillon, and Vincent Balat. Eliom: A core ML language for tierless
952 web programming. In Atsushi Igarashi, editor, *Proceedings of the 14th Asian Symposium*
953 *on Programming Languages and Systems*, APLAS ’16, pages 377–397, Berlin, Heidelberg,
954 November 2016. Springer-Verlag. doi:10.1007/978-3-319-47958-3_20.
- 955 35 Bob Reynders, Frank Piessens, and Dominique Devriese. Gavial: Programming the web with
956 multi-tier FRP. *The Art, Science, and Engineering of Programming*, 4(3):6:1–6:32, February
957 2020. doi:10.22152/programming-journal.org/2020/4/6.
- 958 36 Manuel Serrano, Erick Gallesio, and Florian Loitsch. Hop: a language for programming the
959 web 2.0. In Peri L. Tarr and William R. Cook, editors, *Companion to the 21th Annual ACM*
960 *SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications,*
961 *OOPSLA 2006, October 22-26, 2006, Portland, Oregon, USA*, pages 975–985. ACM, 2006.
962 doi:10.1145/1176617.1176756.
- 963 37 W3C. WS Choreography Description Language. <http://www.w3.org/TR/ws-cdl-10/>, 2004.
- 964 38 Pascal Weisenburger, Mirko Köhler, and Guido Salvaneschi. Distributed system development
965 with ScalaLoc. *Proceedings of the ACM on Programming Languages*, 2(OOPSLA):129:1–129:30,
966 October 2018. doi:10.1145/3276499.
- 967 39 Pascal Weisenburger and Guido Salvaneschi. Multitier modules. In Alastair F. Donaldson,
968 editor, *Proceedings of the 33rd European Conference on Object-Oriented Programming (ECOOP*
969 *’19)*, volume 134 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 3:1–3:29,
970 Dagstuhl, Germany, 2019. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. URL: <http://drops.dagstuhl.de/opus/volltexte/2019/10795>, doi:10.4230/LIPIcs.ECOOP.2019.3.
- 971 40 Pascal Weisenburger, Johannes Wirth, and Guido Salvaneschi. A survey of multitier program-
972 ming. *ACM Computing Surveys*, 53(4), September 2020. doi:10.1145/3397495.
- 973