

Multiparty Languages: The Choreographic and Multitier Cases

Saverio Giallorenzo ✉ 🏠 

Università di Bologna, Italy
INRIA, Sophia Antipolis, France

Fabrizio Montesi ✉ 🏠 

University of Southern Denmark, Odense, Denmark

Marco Peressotti ✉ 🏠 

University of Southern Denmark, Odense, Denmark

David Richter ✉ 

Technical University of Darmstadt, Germany

Guido Salvaneschi ✉ 

University of St. Gallen, Switzerland

Pascal Weisenburger ✉ 

University of St. Gallen, Switzerland

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.2021.22

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.

David Richter: German Federal Ministry of Education and Research (BMBF) iBlockchain Project, grant no. 16KIS0902

Guido Salvaneschi: German Research Foundation (DFG) project no. 383964710, LOEWE initiative (Hesse, Germany) within the Software-Factory 4.0 project, and Swiss National Science Foundation (SNSF) project “Multitier Programming above the Clouds”

Pascal Weisenburger: University of St. Gallen, IPF project no. 1031569

Acknowledgements We thank the anonymous reviewers for their useful feedback and comments.



© 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. 22; pp. 22:1–22:28

Leibniz International Proceedings in Informatics



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

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 runtime, for which programmers do not receive adequate help from mainstream programming abstractions and technology [29, 25, 38].

The quest for finding elegant languages and methodologies that can help with concurrent and distributed programming has been a major focus of the research community for decades, including the seminal actor model and calculus of communicating systems [21, 31]. In this work, we are interested in two kinds of languages that have been recently gaining attention: *choreographic languages* [32, 2] and *multitier languages* [48]. Choreographic languages are designed to express multiparty communication protocols, by providing primitives that make interaction manifest. On the other hand, multitier languages allow for programming computation that spans across several tiers of a distributed system, by providing primitives that allow computation to change location of execution.

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

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

We outline our investigation and contributions:

- In Section 2, we give an overview of the essential features of choreographic and multitier languages. We recap the history of the two approaches and identify their key differences, which lie in perspective (objective vs subjective) and in the modelling of communications (manifest vs non-manifest). We also pinpoint the commonality that classifies choreographic and multitier languages as multiparty.
- In Section 3, we present an example use case for both choreographic and multitier programming, which introduces the concrete choreographic and multitier programming languages that we will use in the rest of our development: Choral [19] and ScalaLoci [46].
- In Section 4, we introduce Mini Choral and Mini ScalaLoci, two representative but minimal languages for choreographic and multitier programming, respectively. Mini Choral and Mini ScalaLoci dispense with the features that are not essential parts of their respective paradigms, which allows us to study how the essential differences can be bridged in the next section.
- In Section 5, we define algorithms for translating programs in Mini Choral to programs in Mini ScalaLoci, and vice versa. The translations deal with the changes in perspective and manifestation of communications between the two paradigms. For example, translating a multitier program into a choreographic one requires synthesising a communication protocol that enacts the necessary communications among participants.

Our translations are not just of inspiration to see the connection between the two paradigms (which we leverage in the next section), but also open a window towards the future sharing of theoretical and practical results. An example for each direction: by translating a multitier program into a choreographic one and then using a choreographic compiler to generate executable code, we can know statically the pattern of communications that will be enacted by the executable code (this property is called “Choreography Compliance” [19] or “EndPoint Projection Theorem” [4]); by translating a choreographic program into a multitier one and then using a multitier compiler to generate executable code, we can reuse all the machinery developed by the multitier community to generate code for different technologies (e.g., the code generated for one participant is in JavaScript for a web browser while the code for another might be code runnable on the Java Virtual Machine for a server).

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

2 Background: Choreographic and Multitier Programming Languages

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

2.1 Choreographic Languages


Choreographic languages are inspired by the famous “Alice and Bob” notation, or security protocol notation [34]. The idea is to define how the different participants of a system should communicate (or interact)—which later inspired also message sequence charts and sequence diagrams [24]. Textual and graphical choreographic languages have already been adopted in industry as specification languages in different settings ranging from business processes, e.g., the choreographic language in OMG’s Business Process Model and Notation, to web services, e.g., W3C’s Web Services Choreography Description Language [37, 45].

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

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

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

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



■ **Figure 1** Choreographic Programming.

■ **Figure 2** Multitier Programming.

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

```

1 Alice.userId -> Bob.x : ch;
2 Bob.x -> Charlie.y : ch2;

```

In the paradigm of *choreographic programming* [32], choreographic languages are full-fledged programming languages: developers write the implementation of an entire multiparty system as a choreography, and then a compiler automatically generates an executable program for each participant. This process is depicted in Figure 1. Choreographies resemble play scripts, written from an external point of view, describing the interactions among all participants. We call this view *objective*. Participants, like **Alice** and **Bob**, are typically referred to as *roles* in choreographies, and the procedure that generates the executable program for each role is called projection (or endpoint projection) [5, 14].

The code in Listing 1 is valid code in the Chor language, the first implementation of choreographic programming [32, 5]. Chor targets microservices: given that code (with appropriate boilerplate), Chor would generate executable programs of microservices that implement **Alice**, **Bob**, and **Charlie**. Choreographic programming has been applied to other settings, e.g., information flow [26], parallel algorithms [13], cyber-physical systems [28, 27], runtime adaptation [14], and integration processes [18].

2.2 Multitier Languages

Multitier languages are inspired by one of the ideas proposed with ambient calculi [6]. In this kind of process calculi, terms express the place (the “ambient”) at which computation occurs. Computations that take place at different locations can be nested, which enables describing multiparty systems. It was later shown that the idea can be combined with well-known abstractions, by developing a variation of λ -calculus with locations called Lambda 5 [33]. This solution prompted the development of *multitier languages* [43, 11, 48], which extend existing programming languages with locations. The term multitier comes from the fact that these languages were mostly developed for web programming, where tiers is used to refer to the typical participants of a web system (e.g., client, backend server, and database).

The crux of a multitier language is the capability of hopping from the point of view of a participant to that of another—the multitier language by Serrano et al. is aptly called “Hop” [43]. When hopping from a participant to another, it is possible to move data from the participant that we are leaving to the participant that we are going to—enabling communication. As an example, consider a remote procedure call from a client to server. In a recent incarnation of multitier programming that builds on the Scala language, ScalaLoci [46], this can be written as follows (for simplicity of presentation, we omit library calls that would be necessary to deal with asynchrony):

```

1  def rpc(input: String): String on Client = on[Client] {
2    val result =
3      on[Server].run.capture(input) {
4        expensiveFunction(input)
5      }.asLocal
6    return result
7  }

```

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

Like choreographic programming languages, multitier languages come with a compiler that turns the multiparty view of the system into executable programs. This process is depicted in Figure 2. Given a multitier program, a multitier compiler generates an executable program for each peer type (in the case of Section 2.2, these would be client and server). The procedure for generating code is called splitting. The nested “dialogues” of peers inside the multitier program depict that a multitier program has many viewpoints, switching regularly from the point of view of a peer to that of another. Nevertheless, code is written with the viewpoint of the peer we are currently in. For this reason, we say that multitier programs adopt a nested *subjective* view. Most multitier programming approaches focus on the Web domain [43, 11, 41, 9, 40, 42] with a client tier, a server tier and potentially a database tier. ScalaLoci broadens the scope to support a greater variety of distributed software architectures such as master–worker patterns often used in cluster computing or peer-to-peer models. Multitier programming has been further used to reason about concurrency [36], security [3] and data locality [7], and applied to information flow [10, 49, 20], software defined networks [35] and high-performance computing [44].

2.3 Towards Linking Choreographic to Multitier Languages

The two communities of choreographic and multitier languages have prolifically evolved independently [2, 48]. They adopted different design principles, and they have found different practical applications—most notably service-oriented computing for choreographies and web development for multitier programming. As a result, they have also developed several features independently (we discuss some of the most important ones in Section 6). In

22:6 Multiparty Languages: The Choreographic and Multitier Cases

addition, the two communities have been facing different challenges. For example, multitier programming languages historically tackle the problem of “impedance mismatch”: the necessity of handling data conversions and heterogeneous execution engines in the web (the Google Web Toolkit is a multitier framework that contributes to this research area). Instead, choreographic programming mainly aimed at achieving “choreography compliance”: providing the guarantee that distributed systems communicate as expected and with desirable properties (like liveness).

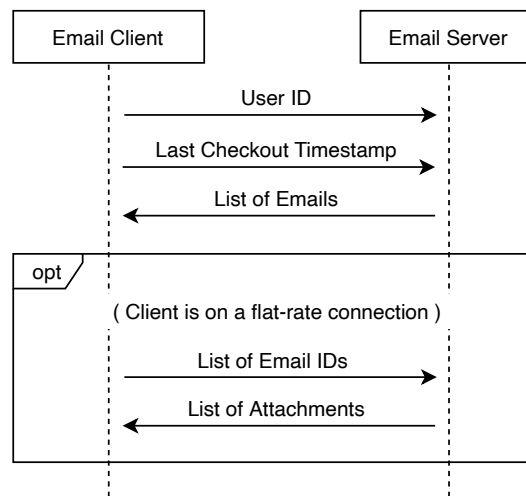
Yet, the two paradigms are clearly linked. We drew Figure 1 and Figure 2 with the intention of highlighting such connection. Indeed, despite differences in both terminologies and methods, the strategies of choreographic and multitier programming languages share a similarity: both define the behaviour 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 [22, 5, 23, 14, 48]. 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* [19]. 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:

```
1 class Example@(Alice, Bob, Charlie) { // the three roles of the protocol
2   DiDataChannel@(Alice,Bob)<Serializable> ch; // channel from Alice to Bob
3   DiDataChannel@(Bob,Charlie)<Serializable> ch2; // channel from Bob to Charlie
4
5   /* constructor omitted */
6
7   public UserID@Charlie run(UserID@Alice userId) { // the protocol
8     UserID@Bob x = ch.<UserID>com(userId); // Alice.userId -> Bob.x : ch
9     return ch2.<UserID>com(x); // Bob.x -> Charlie.y : ch2
10  }
11 }
```

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 `ScalaLoc`. In the next sections, we leverage this common ground and take Choral and `ScalaLoc` as representative languages for their respective paradigms.



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

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

3.2 A Choreographic Programming Implementation with Choral

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

In Choral, objects have types of the form $\tau @ (R_1, \dots, R_n)$, where τ is the interface of the object (as usual), and R_1, \dots, R_n are the roles that collaboratively implement the object. As we see below, Choral supports two notations for denoting the roles over which an object is distributed: the standard form $@(A, \dots, Z)$ and the contracted form $@A$, for objects that belong to one role (shortcut for $@(A)$). Incorporating roles in data types makes distribution manifest at the type level.

■ **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 }

```

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

Considering the description of the implementation of the e-mail fetching protocol, we look at the `updateEmails` method (Line 8). The method does not return a value (`void`) and takes as input the `userId`—which simplifies the user authentication procedure here, for brevity—to identify the user of the `Client` at the `Server`.

In the body, at Line 9, we pass the `userId` to the `Server`. We do this by invoking the method `com` of the `ch SymChannel` giving to it as argument the `userId`. This is done by the expression `userId >> ch::com` which uses the Choral chaining operator `>>` and that corresponds to the expanded expression `ch.com(userId)`. 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)`.

The method `com` of the `SymChannel` transfers the value of the sender given as input into an equivalent representation of the value at the receiver. In this case, the sender is the `Client` (where the `userId` object lives) and the receiver is the `Server`, which stores the result of the communication into variable `id` which is an object of type `UserId` at its location—i.e., `UserId@Server`.

The transfer of the `Timestamp` from the `Client` to the `Server` is similar (Line 10): we retrieve the object from the `clientDB`—invoking method `lastCheckOut`—and we transfer it to the `Server` through the `SymChannel`. Then, to fetch the e-mails, the `Client` receives a transmission from the `Server`. The `Server` interrogates its local database (`serverDB`) by extracting all e-mails

belonging to the `id` of the `Client` and received since its last checkout (indicated by the `timestamp`) and sends them to the `Client` via their shared `SymChannel`. At Line 12, the `Client` uses the received list of `emails` to update its local database (`clientDB`).

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

The `if-else` block at Lines 13–20 allows us to explain the concept of knowledge of choice (a hallmark element of choreographic programming) and how Choral implements it. Briefly, the concept of knowledge of choice indicates a fork in the flow of a program among alternative behaviours, where the concerned roles should coordinate to ensure that they agree on which behaviour they should enact. In choreographic languages, this issue is typically addressed by defining a “selection” primitive to communicate constants drawn from a dedicated set of “labels”, so that the compiler has enough information to build code that can react to choices made by other roles [5, 14]. In Choral, this is implemented by channel methods that can transmit instances of enumerated types between roles. Conveniently, the `SymChannel` used in the example also supports selections via its `select` methods. In Listing 2, we find the implementation of the knowledge of choice of the conditional at Line 14 (where the `Client` “decides” to fetch the attachments) and at Line 20 (which skips the retrieval). In the example, we implement the choice by defining the `Choice` `enum` class at Line 1—note that we use the identifier `Role` for the single role that owns the `Choice` object in its declaration, instantiated at the `Client` at Lines 14 and 20.

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

3.3 A Multitier Programming Implementation with ScalaLoc

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

In ScalaLoc, 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 runtime 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

22:10 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
16    if (ClientLib.isOnFlatRate) {
17      val ids = clientDB.extractIds(emails)
18      clientDB.updateAttachments(
19        on[Server].run.capture(ids) { serverDB.getAttachments(ids) }.asLocal)
20    }
21  }
22 }
```

(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. For simplicity, we use synchronous communication. In general, ScalaLoci allows developers to choose among different *transmitters*, most notably one that wraps local representations of data in futures to account for network delay and communication failures.

The client then uses the received list of `emails` to update its local `clientDB` database (Line 14). Lines 16–20 implement the optional part of the communication logic from Figure 3. If the client is currently using a flat-rate connection—as indicated by the static `ClientLib.isOnFlatRate` method—the client initiates a second server-side computations using `on[Server].run` (Line 19). The client transfers the IDs of the e-mails (retrieved through `extractIds(emails)`)—whose attachments to receive—to the server, which extracts the attachments from its `serverDB` database and returns them to the client, which then updates its local `clientDB` with the received attachments (Line 18).

4 Mini Choreographic and Multitier Languages

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

■ **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 | ::= | $id \mid Exp.id \mid Exp.id(\overline{Exp}) \mid \text{new } id@(\bar{A})(\overline{Exp})$ $\mid \text{lit}(\bar{A}) \mid \text{if} (Exp) \{ Exp \} \text{else} \{ Exp \} \mid Exp; Exp$ $\mid \text{ch_A_B.com}(Exp) \mid \text{ch_A_B.select}(Exp)$ |

4.1 Mini Choral

Listing 4 displays the grammar of Mini Choral. C ranges over class declarations, $Channel$ ranges over channel declarations, $Field$ ranges over class fields, $Method$ ranges over method definitions, $Type$ ranges over type expressions, and Exp ranges over expression terms. The metavariable id ranges over both class names, fields, and variables. We use $\bar{A}, \bar{B}, \bar{C}$ to range over role names. Here and in the remainder of the paper, we use overlines to denote sequences of terms of the same sort and we denote concatenation of sequences using a comma.

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

In Mini Choral, we decided to focus on describing *data flow* and to limit Choral’s expressivity regarding *data distribution*. That is, we allow only the declared **class** to be distributed at multiple roles, while variables belong to only one role, with the exception of *Channels*, which specify the network topology as a set of objects located (and able to transfer single-role objects) between two roles. Specifically, Mini Choral supports only one-way channels (drawn from Choral and called *DiChannels*) of the shape $DiChannel@(\bar{A}, \bar{B}) \text{ ch_A_B}$ —with \bar{A} and \bar{B} roles of the enclosing class. In this work, the loss of expressiveness of the Mini variant with respect to Choral—which supports the definition of multi-role classes/fields without the above limitations—lends itself to simplify the algorithms in our translation in Section 5. In the general case, Choral can express arbitrary channel topologies and user-defined implementations of communications semantics (e.g., asymmetric channels or bidirectional symmetric channels) [19]—whereas most choreographic languages assume a complete topology of channels between all roles in a choreography with a fixed communication semantics [5, 14].

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

Although already captured by the grammar, we include channel invocations of the shape $\text{ch_A_B.com}(Exp)$ and $\text{ch_A_B.select}(Exp)$ to highlight their relevance in the language. *DiChannels* support both methods `com`, meant to transfer data between two roles, and `select`, used to solve knowledge-of-choice challenges in conditionals (that is, informing a role of a local choice made by another role, e.g., by using a conditional) [19]. When using `selects`, we

22:12 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   DiChannel@(Client, Server) ch_Client_Server
3   DiChannel@(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  Unit@Client contextAwareUpdate(List@Client emails) {
18    clientDB.update(emails);
19    if (ClientLib.isOnFlatRate()) {
20      ch_Client_Server.select(Choice@Client.THEN);
21      clientDB.updateAttachments(
22        ch_Server_Client.com(
23          serverDB.getAttachments(
24            ch_Client_Server.com(clientDB.extractIds(emails))))
25      )
26    } else {
27      ch_Client_Server.select(Choice@Client.ELSE); Unit
28    }
29  }
30 }
```

assume that the compiler provides us with a `Choice` `enum` class at one role, with a `THEN` and `ELSE` inhabitants (as presented at Line 1 in Listing 2).

4.1.1 Example: Mini Choral Expressiveness

We conclude the presentation of our minimal choreographic language by illustrating its expressiveness with respect to its reference Choral language with an implementation of the email-fetching protocol presented in Section 3.2, Listing 2.

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

■ **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 ($\overline{id : Type}$) : 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(\overline{id}) { Exp }.asLocal</code> |

4.2 Mini ScalaLoci

Listing 6 displays the grammar of Mini ScalaLoci. L ranges over object declarations, $Peer$ ranges over peer declarations, $Field$ ranges over class fields, $Method$ ranges over method definitions, $Type$ ranges over type expressions, $PlacedType$ ranges over placement type expressions, Exp ranges over expressions, and $PlacedExp$ ranges over placed expressions. The metavariable id ranges over both class names, fields, and variables. We use A , B , C to range over peers.

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

Mini ScalaLoci is able to express different *topologies* rather than being restricted to a *fixed client-server* model. This choice remarks the departure taken by ScalaLoci from other multitier models and implementations [11, 12, 41, 42, 43], which assume a fixed client-server or n -tier architecture of an application. Contrarily, in ScalaLoci, the developer defines an arbitrary number of peers and directional ties between them. In contrast to ScalaLoci, Mini ScalaLoci only supports a single connected peer instance per peer type (drawn from ScalaLoci’s `Single` ties) of the shape `@peer type A <: { type Tie <: Any with Single[A] }`—with A and B peers of the enclosing multitier module. (In Scala, `with` is the operator for constructing intersection types.) In this work, the loss of expressiveness of the Mini variant with respect to ScalaLoci lends itself to simplify the algorithms in our translation in Section 5.

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

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

■ **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) { serverDB.since(id,ts) }.asLocal
14   }
15
16   def contextAwareUpdate(emails: List): Unit on Client = on[Client] {
17     clientDB.update(emails);
18     if (ClientLib.isOnFlatRate()) {
19       updateAttachments(clientDB.extractIds(emails))
20     }
21     else { () }
22   }
23
24   def updateAttachments(emailIds: List): Unit on Client = on[Client] {
25     clientDB.updateAttachments(
26       on[Server].run.capture(emailIds){
27         serverDB.getAttachments(emailIds)
28       }.asLocal
29     )
30   }
31 }

```

4.2.1 Example: Mini ScalaLoci Expressiveness

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

Listing 7 shows the Mini ScalaLoci implementation of the communication scheme in Figure 3. As with Mini Choral, the main notable difference with Listing 3 is that by removing assignments, we rely on method arguments for scoped variable declarations instead. The `updateEmails` method invokes the `getEmails` method, which fetches the emails from the `Server` by sending to it the `id` of the user and the timestamp (`ts`) of the last checkout and transmitting back the result of the extraction on the `serverDB`. Similar to Mini Choral, Mini ScalaLoci also lacks generics, an orthogonal language feature. The lack of generics, however, does not limit the expressiveness of the language to capture the correct topology of the system and communication between the `Server` and the `Client`. After obtaining the emails, we apply method `contextAwareUpdate`, which updates the email database of the client and proceeds to conditionally retrieve the attachments of the fetched emails.

5 Choreographies to Multitier, Multitier to Choreographies

We now define algorithms that translate programs in a Mini language to the other and vice versa. The reason for defining the following algorithms is to present evidence of the existence of a common root at the foundation of the two approaches. We show that the mechanised procedures for their reciprocal translation are relatively simple. In the remainder

■ **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

```

of this section, for brevity, we use the names Choral and ScalaLoci to indicate their Mini counterparts. We first present a translation algorithm from a Choral choreography to a ScalaLoci multitier application (Section 5.1). Afterwards, we show a translation algorithm from a ScalaLoci multitier application to a Choral choreography (Section 5.2).

Perspective translation Multitier and choreographic programming take different perspectives on what parts of the language are annotated with locations. In Choral, all literals are annotated by the role on which they operate, and the location of operators can be inferred by the location of their argument. ScalaLoci assigns peers to expressions, which are then written from the specified peer's perspective.

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

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

Further, we translate between ScalaLoci's way of defining peers and their topology as type members and Choral's way of defining roles as class parameters and their communication channels as class members.

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

The translation algorithms transform **asLocal** in ScalaLoci to an invocation of method **com** of the appropriate channel in Choral and vice versa.

■ **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(ē' ) }"
  | 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(ē') }"
  | 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

```

5.1 From Choreographic Programming to Multitier Programming

Choral choreography classes to ScalaLocI multitier objects Algorithm 1 describes the translation of Choral choreography classes to ScalaLocI multitier objects. We decompose the class definition to be transformed into its identifier *id*, the roles *Role*, the channel declarations *Channel*, the field declarations *Field* and the method definitions *Method*.

Each Choral role definition is translated to a ScalaLocI peer definition. Each channel `DiChannel@A,B` `ch_A_B` between two roles is translated to a single tie, e.g., a directed one-to-one tie, between two peers `@peer type A <: { type Tie <: Any with Single[B] }`.

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

```

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

```

The translation of field definitions from Choral to ScalaLoci is straightforward. In Choral, fields are introduced with a base type and the residing role, followed by the name of the field " $id_{name}@($id_{role}$) id_{type}$ ". In ScalaLoci, fields are introduced as " $\text{val } id_{name}: id_{type} \text{ on } id_{role}$ ". Similarly, method definitions are translated. The algorithm returns a multitier object with the same name and the translated definitions as a body.

Choral choreography expressions to ScalaLoci multitier expressions Algorithm 2 describes the translation of Choral expressions to ScalaLoci: the algorithm matches on the different cases of Choral *Exp* terms and transforms each into the corresponding ScalaLoci code.

For sequencing $e_0; e_1$, both e_0 and e_1 are recursively transformed. If both subexpressions agree on their placement, e.g., $A = B$, the complete sequence is placed on the same peer. More interestingly, if the subexpressions are placed on different peers, we introduce a nested remote block for e'_0 , which executes e'_0 on A and places the overall result of e'_1 on B . For the remote block we generate a capture clause for all method-local variables that are free in e_0 .

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

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

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

■ **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( $\bar{e}'$ )"
  | case "on[A]{ e0.id( $\bar{e}$ ) }" with
    | e'0 ← loci2choral("on[A]{ e0 }")
    | e' ← loci2choral("on[A]{ e }")
    | "e'0.id( $\bar{e}'$ )"
  | 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

```

5.2 From Multitier Programming to Choreographic Programming

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

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

The translation of fields and methods from ScalaLoci to Choral is straightforward. The algorithm returns a Choral class with the same name and the translated definitions as body.

ScalaLoci multitier expressions to Choral choreography expressions Algorithm 4 describes the translation of ScalaLoci expressions to Choral expressions. The algorithm matches on the different cases of ScalaLoci $Expr$ terms and transforms each of them into the corresponding ScalaLoci code.

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

In the case for branching, the translation needs to synthesise select expressions to implement knowledge of choice (recall Section 3.2). Hence, we collect all peers used in the branches and create select statements for all channels between those peers for both branches.

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

| Feature | Choral | ScalaLoci |
|---|--------|-----------|
| Distributed Data Structures (Section 6.1.1) | ✓ | × |
| Dynamic Topologies (Section 6.1.2) | × | ✓ |
| Higher-Order Composition (Section 6.1.3) | ✓ | × |
| Races (Section 6.1.4) | – | ✓ |
| Fault tolerance (Section 6.1.5) | ✓ | ✓ |
| Asynchrony (Section 6.1.6) | ✓ | ✓ |

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

6 A Unified Perspective

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

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

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

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

6.1 Feature Comparison

We now discuss a few features that are important for concurrent and distributed programming. Our discussion is summarised in Table 1, which shows which features are present in Choral and ScalaLoci, respectively (the – in the table means partial support, explained in the relevant paragraph where we discuss the feature). The first four features have evolved separately and give potential for cross-fertilisation, whereas the last two are important features that have been dealt in both worlds (yet separately).

6.1.1 Distributed Data Structures

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

```

1 class BiPair@(A,B)<L@X, R@Y> {
2   private L@A left;
3   private R@B right;
4   public BiPair(L@A left, R@B right) { this.left = left; this.right = right; }
5   public L@A left() { return this.left; }
6   public R@A right() { return this.right; }
7 }

```

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

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

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

```

1 @multitier trait BiPair[L, R] {
2   @peer type A
3   @peer type B
4
5   val left: L on A
6   val right: R on B
7 }

```

While we can define distributed data structures similar to Choral, their usability is more limited: they need to be composed at compile-time, because of `ScalaLoci`'s lack of higher-order composition (see Section 6.1.3).

6.1.2 Dynamic Topologies and Homogenous Behaviours


A feature of `ScalaLoci` that is not covered in its Mini variant is the possibility for peer types to abstract over multiple peer instances of the same type, e.g., a master-worker architecture where a single master can connect to an arbitrary number of homogeneous (i.e., with the same behaviour) worker nodes. Such a feature also enables dynamic topologies where peers can join and leave the system at runtime. A variable number of peer instances is expressed in

■ 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 }

```



ScalaLocis’s peer specification by not using a **Single** tie but a **Multiple** or an **Optional** tie, i.e., an arbitrary number or at most one remote peer of a given type can connect, respectively.

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

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

Deniélou and Yoshida [16] developed a theory for choreographies with homogeneous roles and dynamic topologies by allowing choreographies to be parametrised (also) in collections of roles. Plans for supporting for this feature in Choral are discussed in [19, §7]. In this extension, prefixing a role parameter declaration with *****, as in ***Clients**, specifies that this is a collection of roles. Types are extended with products indexed over collections of role using a syntax similar to Java **for-each** blocks. For instance, the type **forall(Client: Clients) String@Client** represents a “tuple” with a **String** for each role in the collection **Clients**. We can write a scatter-gather channel over a star topology (cf. **MultiClientServer**) as follows:

```

1 abstract class StarChannel@(Server, *Clients) {
2   forall (Client : Clients) { SymChannel@(Server, Client) } star;
3   forall (Client : Clients) { String@Client } scatter(String@Server m);
4   String@Server gather(forall (Client : Clients) { String@Client } ms);
5 }

```

22:22 Multiparty Languages: The Choreographic and Multitier Cases

Method `gather` of `StarChannel` is then translated to ScalaLoc’s primitive `asLocalFromAll` and vice versa. A further extension discussed in [19, §7] is the introduction of existential quantification over roles in role collections. For instance, `with(Client: Clients) {String@Client}` represents a string at some role in the collection `Clients`. We can extend the example above to support any-cast communication as follows:

```
1 abstract class StarChannel@(Server, *Clients) {
2   /* ... */
3   with (Client : Clients) { String@Client } any(String@Server m);
4   String@Server any(with (Client : Clients) { String@Client } m);
5 }
```

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

6.1.3 Higher-Order and First-Class Multiparty Programs

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

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

```
1 class MyClass@(A, B){
2   void passValue(DiChannel@(A, B) ch) {
3     ch.com<Integer>(5@B);
4   }
5 }
```

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

ScalaLoc does not support higher-order composition (no multitier objects as values or dynamic multitier object storage) but at least supports statically-composed modules [47]. The following snippet shows the declaration of a `ClientServer` multitier module that is parameterised over a `Client` and a `Server` peer. The module uses the monitoring functionality provided by the `Monitoring` multitier module, which is parameterised over a `Monitor` and a `Monitored` peer. The `Monitoring` module is instantiated by `mon` inside `ClientServer`. The `ClientServer` module identifies the `Client` peer with the `Monitored` peer and the `Server` peer with the `Monitor` peer and defines their ties accordingly:

```
1 @multitier trait Monitoring {
2   @peer type Monitor { type Tie <: Single[Monitored] }
3   @peer type Monitored { type Tie <: Single[Monitor] }
4 }
5
6 @multitier object ClientServer {
7   @multitier object mon extends Monitoring
8
9   @peer type Client <: mon.Monitored { type Tie <: Single[mon.Monitor] with Single[Server] }
10  @peer type Server <: mon.Monitor { type Tie <: Single[mon.Monitored] with Single[Client] }
11 }
```

Porting higher-order composition from choreographic to multitier languages is an interesting challenge, because the way higher-order values are achieved in the former relies heavily on the objective view of choreographies. Whenever a value is returned in a multitier program, the subjective view of multitier languages requires that the value is located at a single place. It is thus unclear how a higher-order extension of multitier programming should be pursued.

To exemplify the challenge, consider that to return a data structure containing data from two distinct peers A and B, one of the two peers must act as coordinator and collect data from the other, e.g., by nesting `on[A]{ ... on[B]{ ... }.asLocal }`. But this would return a data structure completely located at A, so it does not solve the problem. Alternatively, we could add a multitier operator `par` for running code at different places simultaneously, e.g., `on[A]{ ... } par on[B]{ ... }`. The result of this expression could be a multitier pair containing data at A and B respectively. However, the only way to use this data structure would be to invoke `asLocal` on the two elements of the pair from within an `on[C]` block for some peer C, which would again centralise control.

6.1.4 Races

In this context, by “races” we mean well-behaved and non-deterministic first-come/first-served patterns where two or more roles “race” to communicate with a target role first (and the loser is handled correctly). We distinguish two prototypical scenarios: races among producers and races among consumers.

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

```

1 Future.firstCompletedOf(
2   on[Producer].run { generateValue() }.asLocalFromAll map {
3     case (producerPeerInstance, value) => value map { (producerPeerInstance, _) }
4   })

```

It is not possible to program a race among multiple consumers in ScalaLoci. In general, consumer races represent unexplored territory for the multitier paradigm.

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

```

1 interface ProducerRace@(Producer1, Producer2, Consumer) {
2   Message@Consumer run(Message@Producer1 m1, Message@Producer2 m2);
3 }

```

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

```

1 interface ConsumerRace@(Producer, Consumer1, Consumer2) {
2   BiPair@(Consumer1, Consumer2)<Optional<Message>, Optional<Message>> run(Message@Producer m);
3 }

```

However, the return type of `run` does not guarantee that exactly one consumer receives the message: implementations that deliver the message to both or neither respect the type. As discussed in [19, §7], we can write a precise type if we extend Choral with existential quantification over roles (recall the syntax for existentials at the end of Section 6.1.2) as shown in the example below:

```

1 interface ConsumerRace@(Producer, Consumer1, Consumer2) {
2   with(C : {Consumer1, Consumer2}) { Message@C } run(Message@Producer m);
3 }

```

6.1.5 Fault Tolerance

In ScalaLoci, remote values whose computation or transmission to the local peer instance fail result in a future that is completed with a failure value. Thus, user code can detect a failed remote access and decide how to react appropriately by using library APIs. For example, failed futures can be handled using the typical operators on futures like `recover`:

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

Similarly, Choral does not commit to specific failure handling mechanisms at the language level: programmers can implement their own strategies, e.g., returning errors. An API for channels that is equivalent to the `recover` library method above could look as follows (from the point of view of the caller):

```
chAB.comOrRecover(generateValue(), new OtherValueGenerator@B());
```

where `OtherValueGenerator` has a `run` method equivalent to `generateOtherValue()`. Similar observations hold for timeouts.

ScalaLoci offers some APIs to trigger code when communications with peers in network with dynamic topologies timeout. If dynamic topologies are introduced to Choral, these APIs will become relevant for choreographies as well. We conjecture that they can be imported in a similar way to the one sketched above for recovery.

6.1.6 Asynchrony

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

```
1 interface AsyncDiChannel@(<Sender>, <Receiver>)<T@X> {
2   <S@Y extends T@Y> Future@Receiver<S> com(Promise@Sender<S> v);
3 }
```

6.2 Other Multiparty Languages

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

Software architectures [17, 39] are about organising software systems into well-studied patterns that comprise components and their connections organised in a certain configuration. Architecture description languages (ADL) [30] specify software architectures and the constraints among the architecture components. Different from choreographic and multitier programming, ADLs usually specification languages separate from the implementation. An exception is ArchJava [1] which support specifying a software architecture and enforcing its constraints together with the implementation. Regarding cross-fertilisation, ADLs come equipped with powerful analysis, code synthesis, and runtime-support tools as well as model checkers, which can be also used in multitier and choreographic scenarios to enforce different aspects of correctness.

Partitioned global address space languages (PGAS) [15] are often used in the domain of high-performance computing. The main abstraction is a global memory address space where logical partitions are assigned to processes to maximize data locality. X10 [8] features explicit fork/join operations and provides a sophisticated dependent type system [7] to model the *place* (the heap partition) a reference points to. PGAS languages, similar to multitier and choreographic languages reduce the boundaries between hosts in a distributed system.

7 Conclusion

Choreographic and multitier languages have developed independently, leading to a number of research achievement carried out within two vibrant but separate research communities [2, 32, 48]. In this paper, we discussed the fundamental nature of the programming paradigms based on these languages, isolating the core difference between them. We then showed that, under the cover of syntactic variance, the two approaches are similar enough to be related and to reason about potential cross-fertilisation. Our observations offer a platform for future joint work between the respective communities.

References

- 1 Jonathan Aldrich, Craig Chambers, and David Notkin. ArchJava: Connecting software architecture to implementation. In *Proceedings of the 24th International Conference on Software Engineering, ICSE '02*, pages 187–197, New York, NY, USA, 2002. ACM. doi:10.1145/581339.581365.
- 2 Davide Ancona, Viviana Bono, Mario Bravetti, Joana Campos, Giuseppe Castagna, Pierre-Malo Deniélou, Simon J. Gay, Nils Gesbert, Elena Giachino, Raymond Hu, Einar Broch Johnsen, Francisco Martins, Viviana Mascardi, Fabrizio Montesi, Rumyana Neykova, Nicholas Ng, Luca Padovani, Vasco T. Vasconcelos, and Nobuko Yoshida. Behavioral types in programming languages. *Foundations and Trends in Programming Languages*, 3(2-3):95–230, 2016.
- 3 Ioannis G. Baltopoulos and Andrew D. Gordon. Secure compilation of a multi-tier web language. In *Proceedings of the 4th International Workshop on Types in Language Design and Implementation, TLDI '09*, pages 27–38, New York, NY, USA, 2009. ACM. doi:10.1145/1481861.1481866.
- 4 Marco Carbone, Kohei Honda, and Nobuko Yoshida. Structured communication-centered programming for web services. *ACM Trans. Program. Lang. Syst.*, 34(2):8:1–8:78, 2012. doi:10.1145/2220365.2220367.
- 5 Marco Carbone and Fabrizio Montesi. Deadlock-freedom-by-design: multiparty asynchronous global programming. In Roberto Giacobazzi and Radhia Cousot, editors, *The 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '13, Rome, Italy - January 23 - 25, 2013*, pages 263–274. ACM, 2013. doi:10.1145/2429069.2429101.
- 6 Luca Cardelli and Andrew D. Gordon. Mobile ambients. In Maurice Nivat, editor, *Foundations of Software Science and Computation Structure, First International Conference, FoSSaCS'98, Held as Part of the European Joint Conferences on the Theory and Practice of Software, ETAPS'98, Lisbon, Portugal, March 28 - April 4, 1998, Proceedings*, volume 1378 of *Lecture Notes in Computer Science*, pages 140–155. Springer, 1998. doi:10.1007/BFb0053547.
- 7 Satish Chandra, Vijay Saraswat, Vivek Sarkar, and Rastislav Bodik. Type inference for locality analysis of distributed data structures. In *Proceedings of the 13th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPOPP '08*, pages 11–22, New York, NY, USA, 2008. ACM. doi:10.1145/1345206.1345211.
- 8 Philippe Charles, Christian Grothoff, Vijay Saraswat, Christopher Donawa, Allan Kielstra, Kemal Ebcioglu, Christoph von Praun, and Vivek Sarkar. X10: An object-oriented approach to

- non-uniform cluster computing. In *Proceedings of the 20th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA '05*, pages 519–538, New York, NY, USA, 2005. ACM. doi:10.1145/1094811.1094852.
- 9 Adam Chlipala. Ur/Web: A simple model for programming the web. In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '15*, pages 153–165, New York, NY, USA, 2015. ACM. doi:10.1145/2676726.2677004.
 - 10 Stephen Chong, K. Vikram, and Andrew C. Myers. SIF: Enforcing confidentiality and integrity in web applications. In *Proceedings of 16th USENIX Security Symposium, SS '07*, Berkeley, CA, USA, 2007. USENIX Association. Accessed 2020-05-05. URL: http://usenix.org/events/sec07/tech/full_papers/chong/chong.pdf.
 - 11 Ezra Cooper, Sam Lindley, Philip Wadler, and Jeremy Yallop. Links: Web programming without tiers. In Frank S. de Boer, Marcello M. Bonsangue, Susanne Graf, and Willem P. de Roever, editors, *Formal Methods for Components and Objects, 5th International Symposium, FMCO 2006, Amsterdam, The Netherlands, November 7-10, 2006, Revised Lectures*, volume 4709 of *Lecture Notes in Computer Science*, pages 266–296. Springer, 2006. doi:10.1007/978-3-540-74792-5_12.
 - 12 Ezra E. K. Cooper and Philip Wadler. The RPC calculus. In *Proceedings of the 11th ACM SIGPLAN Conference on Principles and Practice of Declarative Programming, PPDP '09*, pages 231–242, New York, NY, USA, 2009. ACM. doi:10.1145/1599410.1599439.
 - 13 Luís Cruz-Filipe and Fabrizio Montesi. Choreographies in practice. In Elvira Albert and Ivan Lanese, editors, *Formal Techniques for Distributed Objects, Components, and Systems - 36th IFIP WG 6.1 International Conference, FORTE 2016, Held as Part of the 11th International Federated Conference on Distributed Computing Techniques, DisCoTec 2016, Heraklion, Crete, Greece, June 6-9, 2016, Proceedings*, volume 9688 of *Lecture Notes in Computer Science*, pages 114–123. Springer, 2016. doi:10.1007/978-3-319-39570-8_8.
 - 14 Mila Dalla Preda, Maurizio Gabbriellini, Saverio Giallorenzo, Ivan Lanese, and Jacopo Mauro. Dynamic choreographies: Theory and implementation. *Logical Methods in Computer Science*, 13(2), 2017. doi:10.23638/LMCS-13(2:1)2017.
 - 15 Mattias De Wael, Stefan Marr, Bruno De Fraine, Tom Van Cutsem, and Wolfgang De Meuter. Partitioned global address space languages. *ACM Computing Surveys*, 47(4), May 2015. doi:10.1145/2716320.
 - 16 Pierre-Malo Deniérou and Nobuko Yoshida. Dynamic multirole session types. In Thomas Ball and Mooly Sagiv, editors, *Proceedings of the 38th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2011, Austin, TX, USA, January 26-28, 2011*, pages 435–446. ACM, 2011. doi:10.1145/1926385.1926435.
 - 17 David Garlan and Mary Shaw. An introduction to software architecture. Technical report, Pittsburgh, PA, USA, 1994. Accessed 2020-05-05. URL: http://www.cs.cmu.edu/afs/cs/project/vit/ftp/pdf/intro_softarch.pdf.
 - 18 Saverio Giallorenzo, Ivan Lanese, and Daniel Russo. Chip: A choreographic integration process. In *On the Move to Meaningful Internet Systems. OTM 2018 Conferences - Confederated International Conferences: CoopIS, C&TC, and ODBASE 2018, Valletta, Malta, October 22-26, 2018, Proceedings, Part II*, pages 22–40. Springer, 2018. doi:10.1007/978-3-030-02671-4_2.
 - 19 Saverio Giallorenzo, Fabrizio Montesi, and Marco Peressotti. Choreographies as objects. *CoRR*, abs/2005.09520, 2020. URL: <https://arxiv.org/abs/2005.09520>, arXiv:2005.09520.
 - 20 Arjun Guha, Jean-Baptiste Jeannin, Rachit Nigam, Jane Tangen, and Rian Shambaugh. Fission: Secure dynamic code-splitting for JavaScript. In Benjamin S. Lerner, Rastislav Bodík, and Shriram Krishnamurthi, editors, *Proceedings of the 2nd Summit on Advances in Programming Languages (SNAPL 2017)*, volume 71 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 5:1–5:13, Dagstuhl, Germany, 2017. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.SNAPL.2017.5.
 - 21 Carl Hewitt, Peter Bishop, and Richard Steiger. A universal modular Actor formalism for artificial intelligence. In *Proceedings of the 3rd International Joint Conference on Artificial*

- Intelligence*, IJCAI '73, pages 235–245, San Francisco, CA, USA, 1973. Morgan Kaufmann Publishers Inc. Accessed 2020-05-05. URL: <http://ijcai.org/Proceedings/73/Papers/027B.pdf>.
- 22 Kohei Honda, Aybek Mukhamedov, Gary Brown, Tzu-Chun Chen, and Nobuko Yoshida. Scribbling interactions with a formal foundation. In Raja Natarajan and Adegboyega K. Ojo, editors, *Distributed Computing and Internet Technology - 7th International Conference, ICDCIT 2011, Bhubaneshwar, India, February 9-12, 2011. Proceedings*, volume 6536 of *Lecture Notes in Computer Science*, pages 55–75. Springer, 2011. doi:10.1007/978-3-642-19056-8_4.
 - 23 Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty asynchronous session types. *J. ACM*, 63(1):9, 2016. Also: POPL, pages 273–284, 2008. URL: <http://doi.acm.org/10.1145/2827695>, doi:10.1145/2827695.
 - 24 Intl. Telecommunication Union. Recommendation Z.120: Message Sequence Chart, 1996.
 - 25 Tanakorn Leesatapornwongsa, Jeffrey F. Lukman, Shan Lu, and Haryadi S. Gunawi. TaxDC: A taxonomy of non-deterministic concurrency bugs in datacenter distributed systems. In *Proc. of ASPLOS*, pages 517–530, 2016.
 - 26 Alberto Lluch-Lafuente, Flemming Nielson, and Hanne Riis Nielson. Discretionary information flow control for interaction-oriented specifications. In *Logic, Rewriting, and Concurrency*, volume 9200 of *Lecture Notes in Computer Science*, pages 427–450. Springer, 2015.
 - 27 Hugo A. López and Kai Heussen. Choreographing cyber-physical distributed control systems for the energy sector. In *SAC*, pages 437–443. ACM, 2017.
 - 28 Hugo A. López, Flemming Nielson, and Hanne Riis Nielson. Enforcing availability in failure-aware communicating systems. In *FORTE*, volume 9688 of *Lecture Notes in Computer Science*, pages 195–211. Springer, 2016.
 - 29 Shan Lu, Soyeon Park, Eunsoo Seo, and Yuanyuan Zhou. Learning from mistakes: a comprehensive study on real world concurrency bug characteristics. In *Proc. of ASPLOS*, pages 329–339, 2008.
 - 30 Nenad Medvidovic and Richard N. Taylor. A classification and comparison framework for software architecture description languages. *IEEE Transactions on Software Engineering*, 26(1):70–93, January 2000. doi:10.1109/32.825767.
 - 31 Robin Milner. *A Calculus of Communicating Systems*, volume 92 of *Lecture Notes in Computer Science*. Springer, 1980. doi:10.1007/3-540-10235-3.
 - 32 Fabrizio Montesi. *Choreographic Programming*. Ph.D. Thesis, IT University of Copenhagen, 2013. http://www.fabriziomontesi.com/files/choreographic_programming.pdf.
 - 33 Tom Murphy VII, Karl Crary, Robert Harper, and Frank Pfenning. A symmetric modal lambda calculus for distributed computing. In *19th IEEE Symposium on Logic in Computer Science (LICS 2004), 14-17 July 2004, Turku, Finland, Proceedings*, pages 286–295. IEEE Computer Society, 2004. doi:10.1109/LICS.2004.1319623.
 - 34 Roger M. Needham and Michael D. Schroeder. Using encryption for authentication in large networks of computers. *Commun. ACM*, 21(12):993–999, 1978.
 - 35 Tim Nelson, Andrew D. Ferguson, Michael J. G. Scheer, and Shriram Krishnamurthi. Tierless programming and reasoning for software-defined networks. In *Proceedings of the 11th USENIX Conference on Networked Systems Design and Implementation, NSDI '14*, pages 519–531, Berkeley, CA, USA, 2014. USENIX Association. Accessed 2020-05-05. URL: <http://usenix.org/system/files/conference/nsdi14/nsdi14-paper-nelson.pdf>.
 - 36 Matthias Neubauer and Peter Thiemann. From sequential programs to multi-tier applications by program transformation. In *Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '05, pages 221–232, New York, NY, USA, 2005. ACM. doi:10.1145/1040305.1040324.
 - 37 Object Management Group. Business Process Model and Notation. <http://www.omg.org/spec/BPMN/2.0/>, 2011.

- 38 Peter W. O'Hearn. Experience developing and deploying concurrency analysis at facebook. In Andreas Podelski, editor, *Static Analysis - 25th International Symposium, SAS 2018, Freiburg, Germany, August 29-31, 2018, Proceedings*, volume 11002 of *Lecture Notes in Computer Science*, pages 56–70. Springer, 2018. doi:10.1007/978-3-319-99725-4_5.
- 39 Dewayne E. Perry and Alexander L. Wolf. Foundations for the study of software architecture. *ACM SIGSOFT Software Engineering Notes*, 17(4):40–52, October 1992. doi:10.1145/141874.141884.
- 40 Laure Philips, Joeri De Koster, Wolfgang De Meuter, and Coen De Roover. Search-based tier assignment for optimising offline availability in multi-tier web applications. *The Art, Science, and Engineering of Programming*, 2(2):3:1–3:29, December 2018. doi:10.22152/programming-journal.org/2018/2/3.
- 41 Gabriel Radanne, Jérôme Vouillon, and Vincent Balat. Eliom: A core ML language for tierless web programming. In Atsushi Igarashi, editor, *Proceedings of the 14th Asian Symposium on Programming Languages and Systems, APLAS '16*, pages 377–397, Berlin, Heidelberg, November 2016. Springer-Verlag. doi:10.1007/978-3-319-47958-3_20.
- 42 Bob Reynders, Frank Piessens, and Dominique Devriese. Gavial: Programming the web with multi-tier FRP. *The Art, Science, and Engineering of Programming*, 4(3):6:1–6:32, February 2020. doi:10.22152/programming-journal.org/2020/4/6.
- 43 Manuel Serrano, Erick Gallesio, and Florian Loitsch. Hop: a language for programming the web 2.0. In Peri L. Tarr and William R. Cook, editors, *Companion to the 21th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2006, October 22-26, 2006, Portland, Oregon, USA*, pages 975–985. ACM, 2006. doi:10.1145/1176617.1176756.
- 44 Daniel Sokolowski, Jan-Patrick Lehr, Christian Bischof, and Guido Salvaneschi. Leveraging hybrid cloud hpc with multitier reactive programming. In *2020 IEEE/ACM International Workshop on Interoperability of Supercomputing and Cloud Technologies (SuperCompCloud)*, pages 27–32, 2020. doi:10.1109/SuperCompCloud51944.2020.00010.
- 45 W3C. WS Choreography Description Language. <http://www.w3.org/TR/ws-cdl-10/>, 2004.
- 46 Pascal Weisenburger, Mirko Köhler, and Guido Salvaneschi. Distributed system development with ScalaLoc. *Proceedings of the ACM on Programming Languages*, 2(OOPSLA):129:1–129:30, October 2018. doi:10.1145/3276499.
- 47 Pascal Weisenburger and Guido Salvaneschi. Multitier modules. In Alastair F. Donaldson, editor, *Proceedings of the 33rd European Conference on Object-Oriented Programming (ECOOP '19)*, volume 134 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 3:1–3:29, 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.
- 48 Pascal Weisenburger, Johannes Wirth, and Guido Salvaneschi. A survey of multitier programming. *ACM Computing Surveys*, 53(4), September 2020. doi:10.1145/3397495.
- 49 Steve Zdancewic, Lantian Zheng, Nathaniel Nystrom, and Andrew C. Myers. Secure program partitioning. *ACM Transactions on Computer Systems*, 20(3):283–328, August 2002. doi:10.1145/566340.566343.