Concurrent Objects à la Carte

Dave Clarke¹, Einar Broch Johnsen², and Olaf Owe²

¹ CWI, Amsterdam, the Netherlands,
dave@cwi.nl

² Dept. of Informatics, University of Oslo, Norway
{einarj,olaf}@ifi.uio.no

Abstract Services are autonomous, self-describing, technology-neutral software units that can be described, published, discovered, and composed into software applications at run-time. Designing software services and composing services in order to form applications or composite services requires abstractions beyond those found in typical object-oriented programming languages. In this paper, we explore a number of the abstractions used in service-oriented computing and related internet- and web-based programming models in the context of Creol, an executable concurrent object-oriented modelling language with active objects and futures; i.e., features capable of expressing and dealing with asynchronous actions. By adding various abstractions to the modelling language, we demonstrate how a concurrent object language may naturally address many of the requirements of service-oriented computing. The study of language extensions in the restricted setting of a small, high-level modelling language, such as Creol, suggests a cheap way of developing new abstractions for emerging application domains. In this paper, we explore abstractions in the context of service-oriented computing, particularly with regard to dynamic aspects such as service discovery and structuring mechanisms such as groups.

1 Introduction

Service-oriented computing is an emerging computational model in which services are autonomous, self-describing, technology-neutral software units that can be described, published, discovered, and composed into software applications at run-time. A service provides specific functionality to clients in its external environment. In particular, a client may detect new services or better service providers, and negotiate the use of services at run-time, depending on the quality and cost of the services currently available in the environment. This makes service-oriented computing an attractive model for distributed applications such as Web services, e-business, and e-government, but also for ad-hoc, self-configuring, and loosely coupled networks, such as, peer-to-peer networks, sensor-based applications, and mobile systems.

Research on service-oriented computing typically focuses on three aspects: service and data interchangeability (or platform-neutrality), technologies for service detection, and mechanisms for service composition and orchestration. *Interchangeability* is achieved through XML related technologies which support the

exchange of structured or semi-structured data through a shared data model, as well as the publication and description of services. *Service discovery* (or open-endedness) is typically achieved through events and event-handling architectures, and via standards such as UDDI and service repositories. Industrial proposals for the description of *orchestration* mechanisms include BPML, WSFL, XLANG, and BPEL. These languages typically combine communication primitives with work-flow constructs in order to describe the distributed flow of control in flexible ways. Service-oriented computing has been formally studied through a plethora of process algebras (e.g., [8, 9, 23]), which mostly address aspects of orchestration. However, two approaches to service discovery in this context have recently been developed, extending process algebras with Linda tuple spaces to represent service repositories [9] and with semantic subtyping of channels [11].

Object orientation has been criticized for its apparent mismatch with service-oriented computing. This is partly due to a lack of advanced data types, which makes new solutions necessary to elegantly handle XML documents [6], and partly due to its traditionally restrictive communication and concurrency abstractions. However, the venue of truly parallel systems, on distributed as well as multicore platforms, challenges the multithread concurrency model which is predominant in object-oriented systems today and provided by Java and C#. This has led to recent interest in concurrent objects and more flexible object interaction mechanisms [5, 10, 13, 17]. Concurrent objects are reminiscent of the concurrency model of Actors [2] and Erlang [3] and seem promising for deployment on multicore and distributed platforms. Interestingly, concurrent objects have many of the properties described above for service-oriented computing. In particular, their interaction is based on asynchronous communication and does not assume a tight coupling between caller and callee. This leads to a natural notion of asynchronous method calls based on so-called futures [4, 21, 24, 28].

Creol is a concurrent object-oriented language which combines asynchronous method calls with controlled process release [17]. This decouples concurrency, communication, and synchronization in a very flexible way, while supporting a simple proof theory [13] compared to the proof theory of multithread concurrency [1]. A strict encapsulation of fields is enforced by using interfaces for typing objects; these interfaces include a so-called *cointerface* for callback between interacting objects. In this paper, we argue that this flexibility makes it easy to express orchestration patterns. Furthermore, we extend Creol with high-level primitives for service publishing and discovery, and for delegation and a hierarchical notion of service discovery based on groups. We give a type system and formal semantics for this extended language. Finally, we argue through a series of examples for its suitability for service-oriented computing.

The paper is structured as follows. Section 2 introduces Creol with primitives for service-oriented computing, its syntax and type system. Section 3 gives a reduction semantics for the language, in the style of Felleisen and Hieb [15]. Section 4 presents service-oriented computing flavored examples. Section 5 introduces abstractions for groups and their use through another example. Section 6 discusses our approach and some possible extensions in relation to existing work.

2 Service-Oriented Concurrent Objects

In this section, we present an extension of the concurrent object-oriented language Creol to address dynamic service publication and discovery. A concurrent object represents a separate computational unit, conceptually encapsulating its own processor. In Creol, method calls are asynchronous and object variables are typed by interfaces. We use different interfaces to represent the different services offered by a concurrent object. In particular, an interface may require that a cointerface is offered by the service demander.

2.1 Syntax

The language syntax is given in Fig. 1. A program P is a list of interface and class definitions, followed by a method body corresponding to the main method. An interface D may inherit other interfaces and specify a set of method signatures. An interface is a subtype of all the interfaces it inherits. A class L may implement a set of interfaces and provide local fields and method definitions. For simplicity, we ignore class inheritance in this paper. We emphasize the differences with Java. As usual, a method signature declares the return type of the method, its name, and the types of its formal parameters. In addition, the method signature specifies the *cointerface* for the method: i.e., the required type of the object calling the method.

Expressions e are standard apart from the asynchronous method call $e!m(\bar{e})$, the (blocking) read operation $e.\mathbf{get}$, the asynchronous bind operation $\mathbf{bind} \ I$, and the service announcement operation $\mathbf{announce} \ I$. *Statements* s are standard apart from release points $\mathbf{await} \ g$ and $\mathbf{release}$, the delegation statement $\mathbf{forward} \ e!m(\bar{e})$ and the service removal statement $\mathbf{retract} \ e$. *Guards* g are conjunctions of Boolean expressions bv and polling operations $v?$ on futures v . When the guard in an $\mathbf{await}$ statement evaluates to **false**, a process is released, otherwise computation proceeds in the present process. A **release** statement suspends the active process and another suspended process may be rescheduled. Non-deterministic choice $s \sqcap s'$ allows either branch to be selected. The branches of a merge $s \parallel s'$ are interleaved at release points, influencing the flow of control within a process without allowing other processes to execute, unless neither branch is enabled. In addition, the intermediate statement $s \parallel\!\!\! / s'$ appears during reduction; it corresponds to the activation of statement s in the merge of statements s and s' , where statement s' is delayed.

The sequence $t := e!m(\bar{e}'); \mathbf{await} \ t?; v := t.\mathbf{get}$, where t is a future variable, corresponds to a non-blocking method call, while $t := e!m(\bar{e}'); v := t.\mathbf{get}$ (or simply $v := e!m(\bar{e}').\mathbf{get}$) corresponds to a blocking method call.

2.2 Typing

The typing rules are given in Figure 2. Let Γ be a typing context which binds names to types (e.g., $x : T$) and write $\Gamma(x)$ for the type bound to x in Γ , which

$$\begin{array}{ll}
P ::= \overline{D} \ \overline{L} \ \{\overline{T} \ x; s\} & D ::= \text{interface } I \text{ extends } \overline{I} \ \{\overline{M}_s\} \\
M_s ::= T \ m \ (\overline{T} \ x) \ \text{with } I & L ::= \text{class } C \ \text{implements } \overline{I} \ \{\overline{T} \ f; \overline{M}\} \\
M ::= M_s \ \{\overline{T} \ x; s\} & e ::= bv \mid \text{new } C() \mid e.\text{get} \mid e!m(\overline{e}) \mid \text{null} \\
v ::= f \mid x & \mid \text{bind } I \mid \text{announce } I \\
b ::= \text{true} \mid \text{false} & s ::= v := e \mid \text{await } g \mid \text{if } bv \text{ then } s \ \text{fi} \mid \text{skip} \\
bv ::= b \mid v & \mid \text{release} \mid \text{forward } e!m(\overline{e}) \mid \text{retract } e \\
T ::= I \mid \text{bool} \mid \text{fut}(T) & \mid s; s \mid s \square s \mid s \parallel s \mid s \parallel\!\!\! \parallel s \mid \text{return } e \\
T^+ ::= T \mid \text{ad} & g ::= bv \mid v? \mid g \wedge g
\end{array}$$

Figure 1. The language syntax. Variables v are fields (f) or local variables (x), C is a class name, and I an interface name. Lists are indicated by overlining, as in $\overline{e}$. The type language is restricted to booleans, interfaces and futures. (Data types may be added.)

may be undefined. For simplicity, we elide the checking of interface declarations, and assume that all methods declared in an interface have distinct names and that interface extensions are non-circular. Futures have explicit types; i.e., **fut**(T) is the type for futures holding values of type T . In addition, the type **ad** is used to denote the type of service announcements and **ok** the type of well-typed statements. Furthermore, if a class C implements interfaces $\overline{I}$, we conventionally use C as the name for the type **interface** C **extends** $\overline{I}$ $\{ \}$. The subtype relation $I \preceq J$ is induced by the reflexive and transitive closure of the interface extension relation. Data types as well as **ad** are only subtypes of themselves, i.e., **bool** $\preceq$ **bool**, **ad** $\preceq$ **ad**. In addition, if $T \preceq T'$ then **fut**(T) $\preceq$ **fut**(T').

In the typing rules for *expressions*, **new** C gets type C , **null** can get any interface type, and **announce** I gets type **ad**. We assume given an implicit table for interface declarations, so the auxiliary function $\text{lookup}(I, m, \overline{T})$ gives us the return and cointerface types of method m with formal parameters of types $\overline{T}$ in interface I , if such an m is declared in I . The typing of an asynchronous method call $e!m(\overline{e})$, where e has type I and $\overline{e}$ have types $\overline{T}$, can then be explained as follows: If method m is declared in the interface I of the callee e , with return type T and cointerface type J , then asynchronous operation $e!m(\overline{e})$ is typed by **fut**(T) if the caller **this** can provide the required cointerface. Similarly, the asynchronous operation **bind** I is typed by **fut**(I). Similarly, $e.\text{get}$ is of type T if e is a future of type **fut**(T). In the typing of *guards*, the polling of a field f gets type **bool** provided that f is a future. For the typing of *statements*, note that the delegation statement **forward** $e!m(\overline{e})$ is well-typed if the caller of the current method could have called $e!m(\overline{e})$ instead of the call to the current method. It follows in typing rule (FORWARD) that **forward** $e!m(\overline{e})$ is well-typed provided that m has an adequate return type and that the current caller has a type which satisfies the cointerface of m . The service revocation statement **retract** v is well-typed if v has type **ad**. Note that rule (ASSIGN) allows expressions of type T^+ (in contrast to, e.g., (CALL)). The typing of the remaining statements is standard. For the typing of methods, note that the typing context is extended with the two variables **destiny**, which provides the type of the method call's associated future, and **caller**, which is typed by the method's cointerface and

$$\begin{array}{c}
\text{(VAR)} \quad \frac{}{\Gamma \vdash v : \Gamma(v)} \quad \text{(NEW)} \quad \frac{}{\Gamma \vdash \mathbf{new} \, C : C} \quad \text{(NULL)} \quad \frac{}{\Gamma \vdash \mathbf{null} : I} \quad \text{(GET)} \quad \frac{\Gamma \vdash e : \mathbf{fut}(T)}{\Gamma \vdash e.\mathbf{get} : T} \quad \text{(AWAIT)} \quad \frac{\Gamma \vdash g : \mathbf{bool}}{\Gamma \vdash \mathbf{await} \, g : \mathbf{ok}} \\
\\
\text{(CALL)} \quad \frac{\Gamma \vdash e : I \quad \Gamma \vdash \bar{e} : \bar{T} \quad \Gamma(\mathbf{this}) \preceq J \quad \text{lookup}(I, m, \bar{T}) = \langle T, J \rangle}{\Gamma \vdash e!m(\bar{e}) : \mathbf{fut}(T)} \quad \text{(ASSIGN)} \quad \frac{T^+ \preceq \Gamma(v) \quad \Gamma \vdash e : T^+}{\Gamma \vdash v := e : \mathbf{ok}} \quad \text{(SUBSUMPTION)} \quad \frac{\Gamma \vdash e : T \quad T \preceq T'}{\Gamma \vdash e : T'} \\
\\
\text{(ANNOUNCE)} \quad \frac{\Gamma(\mathbf{this}) \preceq I}{\Gamma \vdash \mathbf{announce} \, I : \mathbf{ad}} \quad \text{(RELEASE)} \quad \frac{}{\Gamma \vdash \mathbf{release} : \mathbf{ok}} \quad \text{(COND)} \quad \frac{\Gamma \vdash bv : \mathbf{bool} \quad \Gamma \vdash s : \mathbf{ok}}{\Gamma \vdash \mathbf{if} \, bv \, \mathbf{then} \, s \, \mathbf{fi} : \mathbf{ok}} \\
\\
\text{(FORWARD)} \quad \frac{\Gamma \vdash e : I \quad \text{lookup}(I, m, \bar{T}) = \langle T, J \rangle \quad \Gamma \vdash \bar{e} : \bar{T} \quad \Gamma(\mathbf{caller}) \preceq J \quad T \preceq \text{returnType}(\Gamma)}{\Gamma \vdash \mathbf{forward} \, e!m(\bar{e}) : \mathbf{ok}} \quad \text{(SKIP)} \quad \frac{}{\Gamma \vdash \mathbf{skip} : \mathbf{ok}} \quad \text{(BOOL)} \quad \frac{}{\Gamma \vdash b : \mathbf{bool}} \\
\\
\text{(RETRACT)} \quad \frac{\Gamma \vdash e : \mathbf{ad}}{\Gamma \vdash \mathbf{retract} \, e : \mathbf{ok}} \quad \text{(CHOICE)} \quad \frac{\Gamma \vdash s : \mathbf{ok} \quad \Gamma \vdash s' : \mathbf{ok}}{\Gamma \vdash s \sqcap s' : \mathbf{ok}} \quad \text{(MERGE)} \quad \frac{\Gamma \vdash s : \mathbf{ok} \quad \Gamma \vdash s' : \mathbf{ok}}{\Gamma \vdash s \parallel s' : \mathbf{ok}} \\
\\
\text{(BIND)} \quad \frac{}{\Gamma \vdash \mathbf{bind} \, I : \mathbf{fut}(I)} \quad \text{(POLL)} \quad \frac{\Gamma \vdash v : \mathbf{fut}(T)}{\Gamma \vdash v? : \mathbf{bool}} \quad \text{(G-CONJ)} \quad \frac{\Gamma \vdash g : \mathbf{bool} \quad \Gamma \vdash g' : \mathbf{bool}}{\Gamma \vdash g \wedge g' : \mathbf{bool}} \\
\\
\text{(METHOD)} \quad \frac{\Gamma, \mathbf{destiny} : \mathbf{fut}(T), \mathbf{caller} : I, \bar{x}_1 : \bar{T}_1, \bar{x}_2 : \bar{T}_2 \vdash s : \mathbf{ok}}{\Gamma \vdash T \, m \, (\bar{T}_1 \, x_1) \, \mathbf{with} \, I\{T_2 \, x_2; s\} : \mathbf{ok}} \quad \text{(RETURN)} \quad \frac{\Gamma \vdash e \preceq \text{returnType}(\Gamma)}{\Gamma \vdash \mathbf{return} \, e : \mathbf{ok}} \\
\\
\text{(CLASS)} \quad \frac{\text{implements}(C, \bar{I}) \quad \text{for all } M \in \bar{M} \cdot \Gamma, \mathbf{this} : C, \bar{x} : \bar{T} \vdash M : \mathbf{ok}}{\Gamma \vdash \mathbf{class} \, C \, \mathbf{extends} \, C' \, \mathbf{implements} \, \bar{I} \, \{T \, x; \bar{M}\} : \mathbf{ok}} \quad \text{(COMP)} \quad \frac{\Gamma \vdash s : \mathbf{ok} \quad \Gamma \vdash s' : \mathbf{ok}}{\Gamma \vdash s; s' : \mathbf{ok}}
\end{array}$$

Figure 2. The typing rules. The predicate $\text{implements}(C, \bar{I})$ compares the method signatures of C with those of each interface I in $\bar{I}$ with respect to co- and contravariance requirements for formal parameters and cointerfaces. Function $\text{returnType}(\Gamma) = T$ whenever $\Gamma(\mathbf{destiny}) = \mathbf{fut}(T)$.

provides support for callback, in addition to types for formal parameters and locally declared variables. As usual, the typing rule for classes extends the typing environment with a type for **this** and types for the fields declared in the class. In addition, rule (CLASS) includes a check for compatibility with the declared interfaces $\bar{I}$ of the class C by means of an auxiliary predicate $\text{implements}(C, \bar{I})$ which ensures that for all method signatures $T_1 \, m \, (\bar{T}_2 \, x) \, \mathbf{with} \, A$ declared in an interface I or a superinterface of I (for $I \in \bar{I}$), there is a method declaration $T_3 \, m \, (\bar{T}_4 \, x) \, \mathbf{with} \, B \, \{T \, y; s\}$ in C such that $T_3 \preceq T_1$, $\bar{T}_2 \preceq \bar{T}_1$, and $A \preceq B$.

$$\begin{array}{ll}
\text{config} ::= \epsilon \mid \text{object} \mid \text{msg} \mid \text{service} \mid \text{config config} & \text{fds} ::= \overline{f d} \\
\text{object} ::= (\mathbf{obj} \text{ oid}, C, \text{processQ}, \text{fds}, \text{active}) & \text{active} ::= \text{process} \mid \mathbf{idle} \\
\text{processQ} ::= \epsilon \mid \text{process} \mid \text{processQ processQ} & \text{service} ::= (\mathbf{ad} \text{ adid}, I, \text{oid}) \\
\text{msg} ::= (\mathbf{fut} \text{ fid}, \text{cmd}, \text{oid}, \text{mode}, d) & \text{process} ::= (\overline{T x d}, s) \\
d ::= \text{oid} \mid \text{fid} \mid \text{adid} \mid \mathbf{null} \mid b & \text{cmd} ::= \text{oid}!m(\overline{d}) \mid \mathbf{bind} I
\end{array}$$

Figure 3. Syntax for runtime configurations. Here, d denotes data, including Boolean values (b) and identifiers for objects (oid), futures (fid), and service announcements ($adid$). Processes include code and local state, i.e. local variables with types and values.

3 Operational Semantics

The semantics is a small-step reduction relation on *configurations* of objects, messages, and services (see Fig. 3). An *object* has an identifier, a class, a queue of suspended processes, fields, and an active process. The process **idle** indicates that no method is running in the object. We introduce the notion of a command, cmd , to model asynchronous actions. Commands are reduced using a single rule which spawns off a new thread, in effect, to execute the command. Initially, commands include asynchronous method calls and bind requests. A *future* ($\mathbf{fut} \text{ fid}, \text{cmd}, \text{oid}, \text{mode}, d$) captures the state of command: initially *sleeping*, the command may later becomes *active*, and finally, when *completed*, it stores its result in the future. The value $\text{mode} \in \{\mathbf{s}, \mathbf{a}, \mathbf{c}\}$ represents these three states. A bind request does not use the mode **a**. Default values for types are given by a function default (e.g., $\text{default}(I) = \mathbf{null}$, $\text{default}(\mathbf{bool}) = \mathbf{false}$, and $\text{default}(\mathbf{fut}(T)) = \mathbf{null}$). The *initial configuration* of a program $\overline{L} \{ \overline{T x}; s \}$ contains one object ($\mathbf{obj} \text{ } o, \emptyset, \epsilon, (\overline{T x} \text{ default}(T), s)$).

Reduction takes the form of a relation $\text{config} \rightarrow \text{config}'$. Rules apply to partial configurations and may be applied in parallel. This differs from the semantics of object-oriented languages with a global store [16], but is consistent with Creol's [18] executable Maude semantics [12], and allows true concurrency in the distributed setting. The main rules are given in Figures 4, 5, and 6. The context reduction semantics decomposes a statement into a reduction context and a redex, and reduces the redex [15]. *Reduction contexts* are statements S , expressions E , and guards G with a single hole denoted by $\bullet$:

$$\begin{array}{l}
S ::= \bullet \mid v := E \mid S; s \mid S \parallel s \mid \mathbf{if} G \mathbf{then} s \mathbf{fi} \\
\quad \mid \mathbf{return} E \mid \mathbf{forward} E \mid \mathbf{retract} E \mid \mathbf{await} G \\
E ::= \bullet \mid E.\mathbf{get} \mid E!m(\overline{e}) \mid \text{oid}!m(\overline{d}, E, \overline{e}) \\
G ::= \bullet \mid E? \mid G \wedge g \mid b \wedge G
\end{array}$$

Redexes reduce in their respective contexts; i.e., stat-redexes in S , expr-redexes in E , and guard redexes in G . Redexes are defined as follows:

$$\begin{array}{l}
\text{stat-redexes} ::= x := d \mid f := d \mid \mathbf{await} g \mid \mathbf{skip}; s \mid \mathbf{if} b \mathbf{then} s \mathbf{fi} \\
\quad \mid \mathbf{release} \mid \mathbf{forward} \text{oid}!m(\overline{d}) \mid \mathbf{retract} \text{aid} \mid \mathbf{return} d \\
\text{expr-redexes} ::= x \mid f \mid \text{fid}.\mathbf{get} \mid \text{oid}!m(\overline{d}) \mid \mathbf{new} C() \mid \mathbf{bind} I \mid \mathbf{announce} I \\
\text{guard-redexes} ::= \text{fid}? \mid b \wedge g
\end{array}$$

$$\begin{array}{c}
\text{(RED-CMD)} \\
\hline
\text{fid is fresh} \\
\hline
(\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[\text{cmd}]))) \\
\rightarrow (\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[\text{fid}])) (\mathbf{fut} \text{ fid}, \text{cmd}, \text{oid}, \mathbf{s}, \mathbf{null})
\end{array}$$

$$\begin{array}{c}
\text{(RED-GET)} \\
\hline
(\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[\text{fid.get}])) (\mathbf{fut} \text{ fid}, \text{cmd}, \text{oid}', \mathbf{c}, d) \\
\rightarrow (\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[d])) (\mathbf{fut} \text{ fid}, \text{cmd}, \text{oid}', \mathbf{c}, d)
\end{array}$$

$$\begin{array}{c}
\text{(RED-NEW)} \\
\hline
\text{oid' is fresh} \quad \text{fds}' = \text{defaults}(C') \\
\hline
(\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[\mathbf{new} C']))) \\
\rightarrow (\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[\text{oid'}])) (\mathbf{obj} \text{ oid}', C', \epsilon, \text{fds}', \text{proc}(C, \text{run}, \mathbf{null}, \text{oid}', \epsilon))
\end{array}$$

$$\begin{array}{c}
\text{(RED-POLL)} \\
\hline
b = (\text{mode} \equiv \mathbf{c}) \\
\hline
(\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[\text{fid?}])) (\mathbf{fut} \text{ fid}, \text{cmd}, \text{oid}', \text{mode}, d) \\
\rightarrow (\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[b])) (\mathbf{fut} \text{ fid}, \text{cmd}, \text{oid}', \text{mode}, d)
\end{array}$$

$$\begin{array}{c}
\text{(RED-AWAIT)} \\
\hline
(\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[\mathbf{await} g])) \\
\rightarrow (\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[\mathbf{if} g \text{ then skip else release; await } g \text{ fi}]))
\end{array}$$

$$\begin{array}{c}
\text{(RED-RELEASE)} \\
\hline
S[\mathbf{release}] \neq S'[\mathbf{release}; s \parallel s'] \\
\hline
(\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[\mathbf{release}])) \rightarrow (\mathbf{obj} \text{ oid}, C, pq :: (l, S[\mathbf{skip}]), fds, \mathbf{idle})
\end{array}$$

$$\begin{array}{c}
\text{(RED-RESCHEDULE)} \\
\hline
(\mathbf{obj} \text{ oid}, C, pq :: p :: pq', fds, \mathbf{idle}) \rightarrow (\mathbf{obj} \text{ oid}, C, pq :: pq', fds, p)
\end{array}$$

Figure 4. The context reduction semantics (1/4). In the last rule, pq and pq' may be empty. In RED-NEW, the *proc* function gives an activation of *run* (with *oid'* as caller).

Filling the hole of a context S with an expression r is denoted $S[r]$.

Expressions and guards. In (RED-CMD) spawns asynchronous actions by adding a sleeping future to the configuration, returning its identifier to the caller. In (RED-GET), a read on a future variable blocks the active process until the future is in **completed** mode. Blocking does not reschedule a suspended process. Object creation in (RED-NEW) introduces a new instance of a class C into the configuration (with fields collected from C), and the *run* method is called. In (RED-POLL), a future variable is polled to see if the asynchronous action has completed.

Release and rescheduling. Guards determine whether a process should be released. In (RED-AWAIT), a process at a release point proceeds if its guard is true and otherwise releases. When a process is released, its guard is reused to reschedule the process. When an active process is released in (RED-RELEASE) or terminates, it is replaced by the **idle** process, which allows a process from the process queue to be scheduled for execution in (RED-RESCHEDULE). The rule given here is nondeterministic with respect to which process to choose. Active

$$\begin{array}{c}
\text{(RED-METHOD-BIND)} \\
(\mathbf{obj} \text{ oid}, C, pq, fds, p) (\mathbf{fut} \text{ fid}, \text{oid}!m(\bar{d}), \text{oid}', \mathbf{s}, \mathbf{null}) \\
\rightarrow (\mathbf{obj} \text{ oid}, C, pq :: \text{proc}(C, m, \text{fid}, \text{oid}', \bar{d}), fds, p) (\mathbf{fut} \text{ fid}, \text{oid}!m(\bar{d}), \text{oid}', \mathbf{a}, \mathbf{null}) \\
\\
\text{(RED-RETURN)} \\
\frac{l(\mathbf{destiny}) = \text{fid}}{(\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[\mathbf{return} \text{ d}])) (\mathbf{fut} \text{ fid}, \text{oid}''!m(\bar{d}), \text{oid}', \mathbf{a}, \mathbf{null})} \\
\rightarrow (\mathbf{obj} \text{ oid}, C, pq, fds, \mathbf{idle}) (\mathbf{fut} \text{ fid}, \text{oid}''!m(\bar{d}), \text{oid}', \mathbf{c}, \text{d}) \\
\\
\text{(RED-FORWARD)} \\
\frac{l(\mathbf{destiny}) = \text{fid} \quad l(\mathbf{caller}) = \text{oid}'}{(\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[\mathbf{forward} \text{ oid}''!m(\bar{d})])) (\mathbf{fut} \text{ fid}, -, -, -, -)} \\
\rightarrow (\mathbf{obj} \text{ oid}, C, pq, fds, \mathbf{idle}) (\mathbf{fut} \text{ fid}, \text{oid}''!m(\bar{d}), \text{oid}', \mathbf{s}, \mathbf{null}) \\
\\
\text{(RED-ANNOUNCE)} \\
\frac{\text{adid is fresh}}{(\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[\mathbf{announce} \text{ I}]))} \\
\rightarrow (\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[\text{adid}])) (\mathbf{ad} \text{ adid}, I, \text{oid}) \\
\\
\text{(RED-RETRACT)} \\
(\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[\mathbf{retract} \text{ adid}])) (\mathbf{ad} \text{ adid}, I, \text{oid}) \\
\rightarrow (\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[\mathbf{skip}]))) \\
\\
\text{(RED-SERVICE-BIND)} \\
\frac{I' \preceq I}{(\mathbf{fut} \text{ fid}, \mathbf{bind} \text{ I}, o, \mathbf{s}, \mathbf{null}) (\mathbf{ad} \text{ adid}, I', \text{oid})} \\
\rightarrow (\mathbf{fut} \text{ fid}, \mathbf{bind} \text{ I}, o, \mathbf{c}, \text{oid}) (\mathbf{ad} \text{ adid}, I', \text{oid})
\end{array}$$

Figure 5. The context reduction semantics (2/4). The function *proc* returns the process corresponding to method *m* of class *C* (or a superclass) with default values for the local variables, and **destiny**, **caller**, and the formal parameters, initialized with the actual parameter values (here, *fid*, *oid'*, and $\bar{d}$, respectively).

behavior is initiated by a special method, named *run*, which is activated when the object is created. The interleaving of active and reactive behavior is controlled by means of release points.

Method invocation, return, and delegation. A method call results in an activation on the callee's process queue. As the call is asynchronous, there is a delay between the call and its activation, represented by the sleeping mode of a future. After the call (RED-METHOD-BIND) creates a new process reflecting the activation of the called method, in which the **destiny** variable stores the return address to the call's future. Thus when the process terminates, the result is stored by (RED-RETURN) in the future identified by **destiny**. This future changes its mode to **completed** and the active process becomes **idle**. The delegation of a call to a method of another object in (RED-FORWARD) modifies the future associated with the current call; the new method call is inserted and the mode of the future is reset to **s**, but the *fid* and caller values of the previous call are kept.

$$\begin{array}{c}
\text{(RED-CHOICE1)} \\
\frac{\text{enabled}(s, (fds, l), \mu)}{(\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[s \sqcap s'])) \mu \rightarrow (\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[s])) \mu} \\
\\
\text{(RED-MERGE1)} \quad \text{(RED-MERGE2)} \\
\frac{(\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[s \parallel s']))}{\rightarrow (\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[s \parallel s']))} \quad \frac{(\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[s \parallel s']))}{\rightarrow (\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[s' \parallel s]))} \\
\\
\text{(RED-MERGE-SKIP)} \quad \text{(RED-MERGE-RELEASE1)} \\
\frac{(\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[\mathbf{skip} \parallel s]))}{\rightarrow (\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[s]))} \quad \frac{\text{enabled}(s', (fds, l), \mu)}{(\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[\mathbf{release}; s \parallel s'])) \mu \rightarrow (\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[s' \parallel s])) \mu} \\
\\
\text{(RED-MERGE-RELEASE2)} \\
\frac{\neg \text{enabled}(s', (fds, l), \mu)}{(\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[\mathbf{release}; s \parallel s'])) \mu \rightarrow (\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[\mathbf{release}; (s \parallel s')])) \mu} \\
\\
\text{(RED-ASSIGN-LOCAL)} \quad \text{(RED-VAR-LOCAL)} \\
\frac{(\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[x := d]))}{\rightarrow (\mathbf{obj} \text{ oid}, C, pq, fds, (l[x \mapsto d], S[\mathbf{skip}]))} \quad \frac{(\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[x]))}{\rightarrow (\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[l(x)]))} \\
\\
\text{(RED-ASSIGN-FIELD)} \quad \text{(RED-FIELD)} \\
\frac{(\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[f := d]))}{\rightarrow (\mathbf{obj} \text{ oid}, C, pq, fds[f \mapsto d], (l, S[\mathbf{skip}]))} \quad \frac{(\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[f]))}{\rightarrow (\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[fds(f)]))} \\
\\
\text{(RED-COND1)} \quad \text{(RED-SKIP)} \\
\frac{(\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[\mathbf{if true then } s \text{ fi}]))}{\rightarrow (\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[s]))} \quad \frac{(\mathbf{obj} \text{ oid}, C, pq, fds, (l, \mathbf{skip}; s))}{\rightarrow (\mathbf{obj} \text{ oid}, C, pq, fds, (l, s))} \\
\\
\text{(RED-COND2)} \quad \text{(RED-EMPTY)} \\
\frac{(\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[\mathbf{if false then } s \text{ fi}]))}{\rightarrow (\mathbf{obj} \text{ oid}, C, pq, fds, (l, S[\mathbf{skip}]))} \quad \frac{(\mathbf{obj} \text{ oid}, C, pq, fds, (l, \mathbf{skip}))}{\rightarrow (\mathbf{obj} \text{ oid}, C, pq, fds, \mathbf{idle})} \\
\\
\text{(RED-CONTEXT)} \\
\frac{\text{config} \rightarrow \text{config}'}{\text{config config}'' \rightarrow \text{config}' \text{ config}''} \\
\\
\text{(RED-PARALLEL)} \\
\frac{\text{config } \mu \rightarrow \text{config}' \mu' \quad \text{config}'' \mu \rightarrow \text{config}''' \mu'' \quad \text{dom}(\mu) = \text{dom}(\mu') = \text{dom}(\mu'') \quad \text{dom}(\text{config}') \cap \text{dom}(\text{config}''') = \emptyset}{\text{config config}'' \mu \rightarrow \text{config}' \text{ config}'' \mu' \odot \mu''}
\end{array}$$

Figure 6. The context reduction semantics (3/4). Here, μ denotes a configuration of futures. The rule Red-Merge-Release2 assumes that μ includes all futures in the configuration. The notation $\sigma[v \mapsto d]$ is used to update the binding of v in a state σ . The obvious rule for conjunction is omitted. The function $\text{dom}(\text{config})$ extracts the set of identities of the objects, services, and futures in the configuration config .

Service discovery. In (RED-ANNOUNCE), an object makes public that it provides a given service, and in (RED-RETRACT) an object removes a service an-

nouncement. If a service has been announced by some object, the discovery of that service is done by initiating an asynchronous command **bind** I via (RED-CMD). This request is later bound to an object providing the service in (RED-SERVICE-BIND), which stores the identity of an object announcing the service in the future.

Choice, merge, and release. Either branch of a choice or merge may be selected for reduction, captured by (RED-CHOICE1), (RED-CHOICE2), (RED-MERGE1), and (RED-MERGE2). When a branch of a merge statement completes, (RED-MERGE-SKIP) schedules the other branch. If a release occurs inside a merge, the other branch of the merge is the first candidate for rescheduling — rescheduling is local to a process whenever possible. If both branches release, then the process is released. Let σ map fields and local variables to their values. Process release is based on the predicate *enabled* defined on guards, futures, and states which determines whether a guard will not directly release:

$$\begin{aligned} \text{enabled}(b, \sigma, \mu) &= b \\ \text{enabled}(v, \sigma, \mu) &= \text{enabled}(\sigma(v), \sigma, \mu) \\ \text{enabled}(fid?, \sigma, \mu) &= \text{mode} \equiv \mathbf{c}, \text{ where } (fid, _, _, \text{mode}, _) \in \mu \\ \text{enabled}(g \wedge g', \sigma, \mu) &= \text{enabled}(g, \sigma, \mu) \wedge \text{enabled}(g', \sigma, \mu) \end{aligned}$$

The predicate is lifted to statements; $\text{enabled}(\mathbf{await} \ g, \sigma, \mu) = \text{enabled}(g, \sigma, \mu)$ is the crucial case. In (RED-MERGE-RELEASE1), (RED-MERGE-RELEASE2) and (RED-RELEASE), the contexts and redexes do not factor expressions involving **release** uniquely: these may be factored as both $S[\mathbf{release}]$ and $S'[\mathbf{release}; s \parallel s']$. A clause is added to (RED-RELEASE) to ensure that **release**; $s \parallel s'$ is preferred.

Context and parallel reductions. A reduction applies to a subconfiguration by rule (RED-CONTEXT). In (RED-PARALLEL) futures may be shared between parallel reductions. As the futures inspected by one process may be changed by another, they need to be recomposed in a consistent way. This is handled by a function $\mu \odot \mu'$ which collects futures from μ and μ' and resolves conflicting futures with the same *fid*. New futures are located in *config'* and *config'''*.

Sample reduction. The following is a short reduction sequence illustrating how **bind** works, in combination with the **get** operation for accessing a future. The overline indicates the expression/command being evaluated. The underline indicates the context in which this evaluation occurs.

$$\begin{aligned} & \underline{(\mathbf{obj} \ oid, C, pq, fds, (l, (\overline{\mathbf{bind} \ I}).\mathbf{get}))} \ (\mathbf{ad} \ aid, I', oid') \\ \rightarrow & \underline{(\mathbf{obj} \ oid, C, pq, fds, (l, \overline{\mathbf{fid}.\mathbf{get}}))} \ (\mathbf{fut} \ fid, \overline{\mathbf{bind} \ I}, oid, \mathbf{s}, \mathbf{null}) \ (\mathbf{ad} \ aid, I', oid') \\ \rightarrow & \underline{(\mathbf{obj} \ oid, C, pq, fds, (l, \overline{\mathbf{fid}.\mathbf{get}}))} \ (\mathbf{fut} \ fid, \mathbf{bind} \ I, oid, \mathbf{c}, oid') \ (\mathbf{ad} \ aid, I', oid') \\ \rightarrow & (\mathbf{obj} \ oid, C, pq, fds, (l, oid')) \ (\mathbf{fut} \ fid, \mathbf{bind} \ I, oid, \mathbf{c}, oid') \ (\mathbf{ad} \ aid, I', oid') \end{aligned}$$

4 Examples

We first consider how Creol's high-level concurrency structuring primitives can be applied to orchestration patterns in wide-area computing, by adapting some

examples from Misra and Cook [23]. We then show how the proposed constructs for service discovery and delegation may be used in the context of a restaurant example. To simplify the notation in the examples, we allow expressions $o!\mathbf{m}(\bar{e})$ as program statements if m has return type `void` in the interface of o (and omit the return statement in the body of such methods). Furthermore, we let `Any` be a minimal interface implemented by all classes (thus, the supertype of all other interfaces) and omit specifying the cointerface if it is `Any`.

Example 1. Consider services which provide news on request. These may be modeled by a class `NewsSite` offering an interface `News` to the environment with a method `news`, which for a given date returns a document with the news of that day from a “site”. Let `CNN` and `BBC` be two such sites; i.e., two variables typed by the `News` interface. By calling `CNN!news(d)` or `BBC!news(d)` the news for a specified date d will be (asynchronously) downloaded in XML format from that site. (We do not here consider how XML can be integrated as a data type in Creol, see [27].) Another site distributes emails to clients at request. This site may be modeled by a class `EmailSite` which implements an interface `Email` with a method `send(m, a)`, where m is some message content in XML format and a is the address of a `Client` object. Let `NewsService` be an interface for sites combining news from different sites, as defined by the following class:

```
class NewsServiceSite implements NewsService {
  News CNN; Email email;

  void run() { CNN := new News(); email := new Email() }
  bool newsRelay(Date d, Client a) {
    XML v, fut(XML) t;
    t := CNN!news(d); await t?; v := t.get; email!send(v, a);
    return true
  }
}
```

In this class, a method relays news from `CNN` for a given date d to a given client a by invoking the `send` service of the `email` object with the result returned from the call to `CNN!news`, if `CNN` responds. For simplicity, the method returns **true** upon termination. Note that a `NewsServiceSite` object executing a `newsRelay` process is not blocked while waiting for the `CNN` or the `email` object to respond. Instead the process is suspended until the response has arrived. Once the object responds, the process becomes enabled and may resume its execution, receiving the news and eventually forwarding the news by email. If either of these objects never responds the `NewsServiceSite` object may proceed with its other activity, only the suspended process will never become enabled.

Now consider the case where a client wants news from both objects `CNN` and `BBC`. Because there may be significant delays until an object responds, the client desires to have the news from each object relayed whenever it arrives. This is naturally modeled by the merge operator, as in the following method:

```
bool newsRelay2(Date d, Client a) {
  XML v; fut(XML) t1; fut(XML) t2;
```

```

    t1 := CNN!news(d); t2 := BBC!news(d);
    (await t1?; v:= t1.get; email!send(v,a))
    || (await t2?; v:= t2.get; email!send(v,a));
    return true
}

```

The merge operator allows the news pages from BBC and CNN to be forwarded to the email!send service once they are received. If neither service responds, the whole process is suspended and can be activated again when a response is received. By executing the method, at most two emails are sent.

If the first news page available from either CNN or BBC is desired, the non-deterministic choice operator may be used to invoke the email!send service with the result received from either CNN or BBC, as in the following method:

```

bool newsRelay3(Date d, Client a) {
    XML v; fut(XML) t1; fut(XML) t2;
    t1 := CNN!news(d); t2 := BBC!news(d);
    ((await t1?; v:= t1.get) □ (await t2?; v:= t2.get));
    email!send(v,a); return true
}

```

Here news is requested from both news objects, but only the first response will be relayed. The latest arriving response is ignored.

Example 2. At a restaurant, the host of a dinner party is in charge of ordering wine. For this purpose, the host needs to attract the attention of some passing waiter. If busy, the waiter delegates the order to another waiter. The waiter who eventually deals with the order needs to get a more precise specification of the wine by querying the host; e.g., the desired vintage. The waiter then brings the new bottle to the host of the dinner party. We model this scenario by two interfaces *Sommelier* and *Customer* where the latter is the cointerface of the former, thus enabling callback to the customer. We shall conventionally annotate the interface itself with the cointerface information; in this case, the same cointerface applies to all methods declared in that interface.

```

interface Customer with Waiter {
    int whichVintage()
}
interface Sommelier with Customer {
    bool orderWine(string producer)
}

```

The implementation is given in Fig. 7 and demonstrates the dynamic aspects of the proposed primitives for service-oriented computing. We let a class *Host* implement the interface *Customer* and another class *Waiter* implement *Sommelier*. Waiters initially publish their interface, and retract this service if they have too many customers. Hosts dynamically find waiters when they need to make orders. When too busy, a waiter dynamically finds another waiter and forwards the order. Furthermore, the *Waiter* class demonstrates active behavior, represented by management of service announcements given in the run method,

```

class Host implements Customer, HostDuties {
    bool enoughWine;

    void run() { this!entertainGuests()
    }
    void surveyWineBottle() { Waiter waiter; fut(bool) ack;
        await not(enoughWine); waiter := (bind Waiter).get;
        ack := waiter!orderWine("MyFavoriteWine");
        await ack?; enoughWine := true; release; // enjoy the wine
        this!entertainGuests()
    }
    void entertainGuests() {
        enoughWine := false; this!surveyWineBottle()
    }
    int whichVintage() { return 1970 }
}

class Waiter(int maxLoad) implements Sommelier {
    int noOfCustomers;

    void run() { ad service;
        await noOfCustomers < maxLoad;
        service := announce Waiter;
        await noOfCustomers ≥ maxLoad;
        retract service; this!run()
    }
    bool orderWine(string producer) with Customer {
        fut(int) order; int vintage; Sommelier colleague;
        if (noOfCustomers ≥ maxLoad) then
            colleague := bind Waiter.get;
            forward colleague!orderWine(producer)
        else noOfCustomers := noOfCustomers + 1;
            order := caller!whichWine();
            await order?; vintage := order.get; // serve bottle
            noOfCustomers := noOfCustomers - 1; return true fi
    }
}

```

Figure 7. The restaurant example. In class Waiter, the maxLoad attribute is declared as a class parameter, to be instantiated upon object creation.

interleaved with reactive behavior, represented by the orderWine method. The Host class additionally implements an interface HostDuties, which specifies methods entertainGuests and surveyWineBottle. Note that in the orderWine method of the Waiter class, the callback to whichWine is type correct since Customer is the cointerface.

5 Groups

Groups [7, 20, 22] provide a way of structuring advertisements and hence objects into virtual organisations. For example, groups are used in the JXTA peer-to-peer library [19] to structure peers, advertisements, and channels in order to represent the parties involved in (il)legally downloading a particular film. Groups also offer the possibility of providing access control to service discovery, such as restricting members of the group to only those with legal right to the film. For simplicity, we consider only a very rudimentary access control mechanism, whereby an object has to be a member of a group in order to interact with it. Technically, the operations for service discovery and announcement and their semantics (Sections 2 and 3), will be generalised to the scope of particular groups.

Groups may be created and dissolved, and objects may join and leave groups. Finally, advertisement and discovery happens in the context of a group. For uniformity we denote by **glob** the top-level group in which to advertise and search for advertisements by default. Syntactically groups are treated like objects, and operations on groups like asynchronous method calls. In order to enable groups to be advertised, along with objects, we need to extend the notion of announcement. In Section 2, announcements were purely *subjective*; i.e., an object can only announce itself. In contrast *objective* announcements have the form $e!\mathbf{announce}(e', I)$, where the result of expression e' is advertised using interface I —again assuming that e' has type I —into the group resulting from expression e . In order to discover a service advertised in a group, the object must be a member of that group.

Syntax. The syntax introduced in Fig. 1 is expanded as follows. Let **group** be the type of group identifiers, ranged over by gid , and let I include **group** in order to transparently enable the advertising of groups. Let the *expressions* e include the following operations on groups: **new group**, $e!\mathbf{announce}(e', I)$, and $e!\mathbf{bind}(I)$. Thus, the previous advertisement and discovery operations are redundant: subjective, groupless announcements **announce** I may be encoded as **glob!announce(this, I)** and **bind**(I) as **glob!bind(I)**. Finally, let the *statements* s include the following operations on groups: $e!\mathbf{dissolve}$, $e!\mathbf{join}$, and $e!\mathbf{leave}$. All operations involving groups are asynchronous.

Typing. Figure 8 presents the type rules for the new operations, which simply ensure that these operations apply to groups. Rule (ANNOUNCE-IN-GROUP) requires that the announced value conforms with the announced type. Rule (BIND-IN-GROUP) is the obvious extension to the rule (BIND) of Fig. 2.

Runtime Syntax. The runtime syntax introduced in Fig. 3 is expanded as follows. Let id be either *oid* or *gid*. We expand the notion of command as follows:

$$\begin{aligned} cmd ::= & \dots \mid \mathbf{new\ group} \mid gid!\mathbf{dissolve} \mid gid!\mathbf{join} \\ & \mid gid!\mathbf{leave} \mid gid!\mathbf{announce}(id, I) \mid gid!\mathbf{bind}(I) \end{aligned}$$

Configurations *config* are extended with a representation (**gr** $gid, oids, aids$) for a group with identifier gid , whose members are objects represented by the set

$$\begin{array}{c}
\text{(GROUP-CREATE)} \qquad \text{(GROUP-DISSOLVE)} \qquad \text{(GROUP-JOIN)} \\
\frac{}{\Gamma \vdash \mathbf{new\ group} : \mathbf{fut}(\mathbf{group})} \quad \frac{\Gamma \vdash e : \mathbf{group}}{\Gamma \vdash e!\mathbf{dissolve} : \mathbf{ok}} \quad \frac{\Gamma \vdash e : \mathbf{group}}{\Gamma \vdash e!\mathbf{join} : \mathbf{ok}} \\
\text{(GROUP-LEAVE)} \qquad \text{(ANNOUNCE-IN-GROUP)} \qquad \text{(BIND-IN-GROUP)} \\
\frac{\Gamma \vdash e : \mathbf{group}}{\Gamma \vdash e!\mathbf{leave} : \mathbf{ok}} \quad \frac{\Gamma \vdash e : I \quad \Gamma \vdash e' : \mathbf{group}}{\Gamma \vdash e'!\mathbf{announce}(e, I) : \mathbf{fut}(\mathbf{ad})} \quad \frac{\Gamma \vdash e : \mathbf{group}}{\Gamma \vdash e!\mathbf{bind}(I) : \mathbf{fut}(I)}
\end{array}$$

Figure 8. Typing for group-related operations

oids and containing the advertisements *aids*. The reduction contexts become:

$$\begin{aligned}
S ::= & \dots \mid E!\mathbf{dissolve} \mid E!\mathbf{join} \mid E!\mathbf{leave} \\
& \mid E!\mathbf{announce}(e, I) \mid d!\mathbf{announce}(E, I) \mid E!\mathbf{bind}(I)
\end{aligned}$$

The redexes become:

$$\begin{aligned}
\text{stat-redexes} ::= & \dots \mid gid!\mathbf{dissolve} \mid gid!\mathbf{join} \mid gid!\mathbf{leave} \\
\text{expr-redexes} ::= & \dots \mid gid!\mathbf{bind}(I) \mid \mathbf{new\ group} \mid gid!\mathbf{announce}(id, I)
\end{aligned}$$

Reduction Rules. The rule (RED-CMD) in Fig. 5 applies to the asynchronous evaluation of commands. The reduction rules for the new operations are given in Figure 9. As all new operations are asynchronous, the reduction rules operate in the context of some future data structure, generally depending upon the fact that the caller is stored within this data structure in order to perform a primitive access control check. Apart from (RED-CREATE), (RED-JOIN) and (RED-LEAVE) an object needs to be a member of the group in order to interact with it. Rule (RED-CREATE) creates a new group which contains no advertisements or members and (RED-DISSOLVE) removes an entire group. Rule (RED-JOIN) and (RED-LEAVE) allow an object to join and leave a group, by modifying the list of group members. Rule (RED-ANNOUNCE) creates a new announcement and places it in the group and (RED-RETRACT) removes an announcement from a group, generalising the earlier version of this rule. Rule (RED-BIND) binds to an announcement in the scope of a given group.

Example 3. Groups may be used to structure the services of auction houses. The underlying services are managed by the `ItemAuction` class, corresponding to an item for sale in an auction. The items are grouped into auctions, presumably the collection of items to be auctioned during a particular day or perhaps of a specific kind of item (e.g., guitars). Auctions are organised into auction houses, which again are groups.

To simplify the presentation, variable initialization may occur in declarations, with the syntax $T \ v = e$. Furthermore, we use a simple form of multicast; i.e., for each object in $\bar{e}$, $\bar{e}!m(\bar{e}')$ produces one future and $e'!\mathbf{announce}(\bar{e}, I)$ makes announcements, adding the objects to the group e' . Some basic functions on the data type `Set [T]` are assumed to be defined, such as *add* and *remove*.

The following class illustrates how some auctions could be created and populated from an external source.

$$\begin{array}{c}
\text{(RED-CREATE)} \\
\frac{gid \text{ is fresh}}{(\text{fut } fid, \text{new group}, oid, \mathbf{s}, \text{null}) \rightarrow (\text{fut } fid, \text{new group}, oid, \mathbf{c}, gid) (\text{gr } gid, \epsilon, \epsilon)} \\
\\
\text{(RED-DISSOLVE)} \\
\frac{oid \in oids}{(\text{fut } fid, gid!\text{dissolve}, oid, \mathbf{s}, \text{null}) (\text{gr } gid, _, oids) \rightarrow (\text{fut } fid, gid!\text{dissolve}, oid, \mathbf{c}, \text{null})} \\
\\
\text{(RED-JOIN)} \\
(\text{fut } fid, gid!\text{join}, oid, \mathbf{s}, \text{null}) (\text{gr } gid, oids, aids) \rightarrow (\text{fut } fid, gid!\text{join}, oid, \mathbf{c}, \text{null}) (\text{gr } gid, oids \cup \{oid\}, aids) \\
\\
\text{(RED-LEAVE)} \\
(\text{fut } fid, gid!\text{leave}, oid, \mathbf{s}, \text{null}) (\text{gr } gid, oids, aids) \rightarrow (\text{fut } fid, gid!\text{leave}, oid, \mathbf{c}, \text{null}) (\text{gr } gid, oids \setminus \{oid\}, aids) \\
\\
\text{(RED-ANNOUNCE)} \\
\frac{oid \in oids \quad aid \text{ is fresh}}{(\text{fut } fid, gid!\text{announce}(id, I), oid, \mathbf{s}, \text{null}) (\text{gr } gid, oids, aids) \rightarrow (\text{fut } fid, gid!\text{announce}(id, I), oid, \mathbf{c}, aid) (\text{ad } aid, I, id) (\text{gr } gid, oids, aids \cup \{aid\})} \\
\\
\text{(RED-RETRACT)} \\
\frac{oid \in oids \quad aid \in aids}{(\text{fut } fid, \text{retract } aid, oid, \mathbf{s}, \text{null}) (\text{ad } aid, I, id) (\text{gr } gid, oids, aids) \rightarrow (\text{fut } fid, \text{retract } aid, oid, \mathbf{c}, \text{null}) (\text{gr } gid, oids, aids \setminus \{aid\})} \\
\\
\text{(RED-BIND)} \\
\frac{oid \in oids \quad aid \in aids}{(\text{fut } fid, gid!\text{bind}(I), oid, \mathbf{s}, \text{null}) (\text{ad } aid, I, id) (\text{gr } gid, oids, aids) \rightarrow (\text{fut } fid, gid!\text{bind}(I), oid, \mathbf{c}, id) (\text{ad } aid, I, id) (\text{gr } gid, oids, aids)}
\end{array}$$

Figure 9. The context reduction semantics (4/4) for group-related operations.

```

class Main(ItemDatabase db) {
  void run() {
    group ah1 = (new group).get; this!populate("Christie's", ah1);
    group ah2 = (new group).get; this!populate("Dave's", ah2);
    group as3 = (new group).get; this!populate("Einar og Olaf's", ah3);
  }
  void populate(String name, group house) {
    Set[Item] items;
    group auction = (new group).get;
    house!join; house!announce(auction, group);
    // populate auction from external source
    items := db!itemsFor(name).get;
    auction!announce(items, Item) // multicast
  }
}

```


The class `ItemAuction` manages the auction of an individual item, including a set of the bidders for that item, and notifications to those bidders when the auction status changes. The method `bid` handles individual bids, relying on the mutual exclusion provided by Creol objects, and notifies other bidders. Class `ItemAuction` implements an interface `Item` with `Bidder` as cointerface and methods `showInterest`, `loseInterest`, and `bid`. (The `Bidder` interface and its implementation are given further below.)

```

class ItemAuction implements Item {
  int currentBid = 0;
  Bidder holder = null;
  bool sold = false;
  Set[Bidder] bidders = new Set();

  bool bid(int amount) with Bidder {
    if (sold or amount <= current) then return false fi;
    bidders := add(bidders, caller);
    currentBid := amount; holder := caller;
    remove(bidders, caller)!higherBid(amount); // multicast
    return true
  }
  // allow a bidder to (un)subscribe to an item auction
  void showInterest() with Bidder { bidders := add(bidders, caller) }
  void loseInterest() with Bidder { bidders := remove(bidders, caller) }

  void run() { await currentBid > 0; this!watchBids() }

  void watchBids() {
    int previousBid = currentBid;
    release; release; release; // going once, twice, three times
    if currentBid > previousBid then this!watchBids()
    else sold := true; // accept current bid
      holder!success();
      bidders!auctionClosed() fi // multicast
  }
}

```

Specific items can be implemented by extending `Item`, giving a more specific description (in lieu of meta-data).

```

interface Guitar extends Item {}
interface Gibson extends Guitar {}
interface Firebird extends Gibson {}
interface FirebirdI extends Firebird {}
interface NonReverseFirebird extends Firebird {}

```

`Bidder` is an interface implemented by potential buyers in order to receive notifications from the auctions they follow. The notifications include that a higher bid has been made, that the auction has closed, and that a bid was successful. In each case, the **caller** is the item to which the notifications refer.

```

interface Bidder with Item {
  void higherBid(int bid);
  void auctionClosed();
  void success()
}

```

Buyers use service discovery to find groups denoting auction houses, and within an auction house, an auction will be found, within which the specific items of interest can be found.

```

class Buyer implements Bidder {
  void run() {
    // find an auction house -- needs meta-data
    group ah = glob!bind(group).get; ah!join;
    group auction = ah!bind(group).get; auction!join;
    Item item = auction!bind(Firebird).get;
    item!showInterest()
  }
  // Implementation of the Bidder interface
  void higherBid(int bid) with Item {
    if (bid < myMax) then caller!bid(bid + 10) fi
  }
  void auctionClosed() with Item {...} // quit, go home unhappy
  void success() with Item {...} // quit, go home happy
}

```

6 Discussion

This paper explored a number of abstractions for service-oriented computing and related internet- and web-based programming models in the context of the concurrent object-oriented (modelling) language Creol. In particular, we have considered abstractions of dynamic aspects such as service discovery and structuring mechanisms such as groups. By adding abstractions to the modelling language, we showed how Creol satisfies many of the requirements of service-oriented computing. Directly introducing abstractions into the language and implementing them in Maude [12], is a quick way of enabling model development for emerging application domains, simultaneously providing a rapid way of evaluating alternative design decisions. There are many ways to extend Creol to enhance its ability to accurately model real world systems. We now explore a few of these.

Open system and failure models New services can come on-line at any time. This can be modelled by external interaction with a running Creol expression. External interaction may be captured by a labelled transition—and the existing set of rewrite rules would need to be modified to propagate such transitions:

$$\begin{array}{c}
\text{(META-INJECT)} \\
\epsilon \xrightarrow{\text{inject } object} object
\end{array}$$

The initiation of Creol's class update mechanism [29] may be modelled this way, and may be used to provide runtime upgrade of services.

Dually, it is possible to model when services go off-line, either temporarily or permanently. The following rule considers an explicit **remove** command:

$$\begin{array}{c} \text{(RED-REMOVE)} \\ (\mathbf{fut} \text{ } fid, \mathbf{remove} \text{ } oid, oid', \mathbf{s}, _) (\mathbf{obj} \text{ } oid, C, processQ, fds, active) \\ \rightarrow (\mathbf{fut} \text{ } fid, \mathbf{remove} \text{ } oid, oid', \mathbf{c}, \mathbf{null}) \end{array}$$

A similar meta-rule interacting with the environment could model the operation.

In general, operations in distributed systems can only be considered as partial due to (intermittent) failures of the network. A partial solution in the context of Creol is to modify asynchronous calls to allow the possibility of timing-out. A uniform way to achieve this is to add time-outs to future objects, as follows:

$$(\mathbf{fut} \text{ } fid, cmd, caller, \underline{timer}, mode, result),$$

where *timer* has one of the following values **timer infinity** or **timer** *n*, for *n* ≥ 0. An additional mode **to** is introduced to model the occurrence of a time-out. The new, under-the-hood, reduction rules are:

$$\begin{array}{c} \text{(RED-TIMER)} \\ \mathbf{timer} (n + 1) \rightarrow \mathbf{timer} \text{ } n \end{array} \quad \frac{\begin{array}{c} \text{(RED-TIMEOUT)} \\ status \neq \mathbf{c} \end{array}}{(\mathbf{fut} \text{ } fid, cmd, caller, \mathbf{timer} \text{ } 0, status, result) \rightarrow (\mathbf{fut} \text{ } fid, cmd, caller, \mathbf{timer} \text{ } 0, \mathbf{to}, result)}$$

Changes at the language level, such as the interaction between **get** and **await** and time-outs would need to be done, perhaps via an exception mechanism. Maude has a real-time extension [25] to help implement these features.

Enhancing Groups Although we used subinterfaces (of *Guitar*) to enable more precise service discovery, group discovery was rather weak, as there were no means for distinguishing one group from another. More precise groups can be specified by applying inheritance to groups, as follows:

```
group QualityGroup;
group GuitarGroup;
group GibsonGroup extends GuitarGroup, QualityGroup;
```

Group discovery would be based on one of the supergroups; e.g., *QualityGroup*. Group creation would then be extended to create a specific kind of group:

```
new GibsonGroup;
```

Further possibilities include limiting the types that can be advertised in a group:

```
group GuitarGroup of Guitar;
```

Additional meta-data such as tags and certificates (for access control) [19] could be supplied to enhance the quality of service discovery.

Additional group operations would certainly increase Creol’s modelling power; e.g., broadcast to a group [20]. Likewise, a stream-based approach to asynchronous responses, as in JXTA [19], to model multiple potential candidate bindings, would enable a client to better select the most appropriate service.

Enhancing the security model underlying our groups would make them more appropriate for modelling real-world applications. Such measures are necessary since one would not generally want everyone to be able to dissolve the group, and some members may be able to advertise, whereas others may not.

Other directions for future work include adding behavioural types (e.g., session types [14]) to place more static control on the use of objects or services, and to add a coordination layer (channels, connectors, or tuple space(s) [26]) between the objects or services, rather than allowing direct communication.

Acknowledgment. We are grateful to Frank de Boer for interesting discussions on service binding in the context of object orientation.

References

1. E. Ábrahám-Mumm, F. S. de Boer, W.-P. de Roever, and M. Steffen. Verification for Java’s reentrant multithreading concept. In *International Conference on Foundations of Software Science and Computation Structures (FOSSACS’02)*, volume 2303 of *Lecture Notes in Computer Science*, pages 5–20. Springer, Apr. 2002.
2. G. A. Agha. *ACTORS: A Model of Concurrent Computations in Distributed Systems*. The MIT Press, Cambridge, Mass., 1986.
3. J. Armstrong. *Programming Erlang: Software for a Concurrent World*. Pragmatic Bookshelf, 2007.
4. H. G. Baker and C. E. Hewitt. The incremental garbage collection of processes. *ACM SIGPLAN Notices*, 12(8):55–59, 1977.
5. N. Benton, L. Cardelli, and C. Fournet. Modern concurrency abstractions for C[#]. *ACM Transactions on Programming Languages and Systems*, 26(5):769–804, Sept. 2004.
6. G. M. Bierman, E. Meijer, and W. Schulte. The essence of data access in C_ω. In A. P. Black, editor, *Proc. 19th European Conference on Object-Oriented Programming (ECOOP’05)*, volume 3586 of *Lecture Notes in Computer Science*, pages 287–311. Springer, 2005.
7. K. P. Birman. The process group approach to reliable distributed computing. *Communications of the ACM*, Dec. 1993.
8. M. Boreale, R. Bruni, L. Caires, R. D. Nicola, I. Lanese, M. Loreti, F. Martins, U. Montanari, A. Ravara, D. Sangiorgi, V. T. Vasconcelos, and G. Zavattaro. SCC: A service centered calculus. In M. Bravetti, M. Núñez, and G. Zavattaro, editors, *Proc. 3rd Intl. Workshop on Web Services and Formal Methods*, volume 4184 of *Lecture Notes in Computer Science*, pages 38–57. Springer, Sept. 2006.
9. M. Bravetti and G. Zavattaro. Service oriented computing from a process algebraic perspective. *Journal of Logic and Algebraic Programming*, 70(1):3–14, 2007.
10. D. Caromel and L. Henrio. *A Theory of Distributed Object*. Springer, 2005.
11. G. Castagna, R. D. Nicola, and D. Varacca. Semantic subtyping for the pi-calculus. In *Proc. 20th IEEE Symposium on Logic in Computer Science (LICS’05)*, pages 92–101. IEEE Computer Society Press, 2005.

12. M. Clavel, F. Durán, S. Eker, P. Lincoln, N. Martí-Oliet, J. Meseguer, and J. F. Quesada. Maude: Specification and programming in rewriting logic. *Theoretical Computer Science*, 285:187–243, Aug. 2002.
13. F. S. de Boer, D. Clarke, and E. B. Johnsen. A complete guide to the future. In R. de Nicola, editor, *Proc. 16th European Symposium on Programming (ESOP’07)*, volume 4421 of *Lecture Notes in Computer Science*, pages 316–330. Springer, Mar. 2007.
14. M. Dezani-Ciancaglini, D. Mostrous, N. Yoshida, and S. Drossopoulou. Session types for object-oriented languages. In D. Thomas, editor, *ECOOP’06*, volume 4067 of *Lecture Notes in Computer Science*, pages 328–352. Springer, 2006.
15. M. Felleisen and R. Hieb. The revised report on the syntactic theories of sequential control and state. *Theoretical Computer Science*, 103(2):235–271, 1992.
16. M. Flatt, S. Krishnamurthi, and M. Felleisen. A programmer’s reduction semantics for classes and mixins. In *Formal Syntax and Semantics of Java*, volume 1523 of *Lecture Notes in Computer Science*, pages 241–269. Springer, 1999.
17. E. B. Johnsen and O. Owe. An asynchronous communication model for distributed concurrent objects. *Software and Systems Modeling*, 6(1):35–58, Mar. 2007.
18. E. B. Johnsen, O. Owe, and I. C. Yu. Creol: A type-safe object-oriented model for distributed concurrent systems. *Theoretical Computer Science*, 365(1–2):23–66, Nov. 2006.
19. JXTA Project. <http://jxta.org>.
20. D. Lea. Objects in groups. <http://gee.cs.oswego.edu/dl/groups/groups.html>, December 1993.
21. B. H. Liskov and L. Shrira. Promises: Linguistic support for efficient asynchronous procedure calls in distributed systems. In D. S. Wise, editor, *Proc. SIGPLAN Conference on Programming Language Design and Implementation (PLDI’88)*, pages 260–267. ACM Press, June 1988.
22. S. Maffei. Electra: Making distributed programs object-oriented. In *Proceedings of USENIX Symposium on Experiences with Distributed and Multiprocessor Systems*, Sept. 1993.
23. J. Misra and W. R. Cook. Computation orchestration: A basis for wide-area computing. *Software and Systems Modeling*, 6(1):83–110, Mar. 2007.
24. J. Niehren, J. Schwinghammer, and G. Smolka. A concurrent lambda calculus with futures. *Theoretical Computer Science*, 364:338–356, 2006.
25. P. C. Ölveczky and J. Meseguer. Semantics and pragmatics of Real-Time Maude. *Higher-Order and Symbolic Computation*, 20(1–2):161–196, 2007.
26. G. A. Papadopoulos and F. Arbab. Coordination models and languages. In M. Zelkowitz (Ed.), *The Engineering of Large Systems*, volume 46 of *Advances in Computers*, pages 329–400. Academic Press, 1998.
27. A. Torjusen, O. Owe, and G. Schneider. Towards integration of XML in the Creol object-oriented language. In F. E. Sandnes, editor, *Proc. Norsk Informatikkonferanse 2007 (NIK’07)*, pages 107–111. Tapir Akademisk Forlag, 2007.
28. A. Welc, S. Jagannathan, and A. Hosking. Safe futures for Java. In *Proc. 20th Conf. on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA’05)*, pages 439–453. ACM Press, 2005.
29. I. C. Yu, E. B. Johnsen, and O. Owe. Type-safe runtime class upgrades in Creol. In R. Gorrieri and H. Wehrheim, editors, *Proc. 8th International Conference on Formal Methods for Open Object-Based Distributed Systems (FMOODS’06)*, volume 4037 of *Springer*, pages 202–217. Springer, June 2006.